All in the Family: Bunching Elasticities and Intra-Family Income Shifting*

Anand Bharadwaj † Robert Breunig ‡

February 10, 2026

Abstract

Bunching estimators are widely used in public finance to estimate the elasticity of taxable income. However, their interpretation becomes ambiguous when taxpayers can shift income, as it becomes necessary to distinguish bunching that reflects genuine changes in economic activity from bunching induced by ex-post income shifting. This paper shows that, under general conditions, fundamental bunching at the individual level should imply corresponding bunching in the aggregate taxable income of a "shiftable syndicate." The shiftable syndicate framework accommodates diverse forms of income-shifting, including both intertemporal and intra-family shifting. Consequently, the absence of bunching in aggregate income falsifies the presence of fundamental bunching at the individual level. To test this empirically, we develop new statistical methods for assessing whether an empirical income distribution is well approximated by a globally smooth distribution. Applying these methods to Australian tax data, we find essentially no bunching in total family income, providing strong evidence that much of the observed bunching in individual taxable income arises from intra-family income shifting.

Keywords: Taxation and incentives; bunching; income shifting; density estimation
JEL Classification: H21; H24; H26; D31, C14

*We are very grateful to Alan Auerbach, Muhammad Bashir, Tim Cejka, Mark Gergen, Stephen King, Emi Nakamura, Emmanuel Saez, Joel Slemrod, Chris Shannon, Jón Steinsson and Dustin Swonder for many valuable comments and suggestions. We also wish to thank the staff at the Australian Taxation Office, particularly Kim Cvetanovska, Son Nguyen and Justin Holland, for facilitating data access and for ongoing guidance and advice. Bharadwaj gratefully acknowledges support from the Robert D. Burch Center for Tax Policy and Public Finance.

†Department of Economics, University of California, Berkeley. Email: anand_bharadwaj@berkeley.edu.
‡Crawford School of Public Policy, Australian National University. Email: robert.breunig@anu.edu.au.

1 Introduction

Since the seminal work of Saez (2010), bunching estimators have become a central tool in public finance to quantify the responsiveness of individuals to changes in tax rates or policies. In the canonical application of this methodology, the degree of "bunching" or excess mass around kink points of the tax schedule is in principle informative about the elasticity of taxable income, which in several standard models is a sufficient statistic for the efficiency of the tax system (Feldstein, 1999; Kopczuk, 2025). However, what drives observed bunching is something of an open question, with important consequences for the interpretation of bunching elasticities. Is bunching driven by fundamental responses in earned income or is it merely symptomatic, as Gordon and Slemrod (1998) provocatively suggest, of income shifting between different tax bases?

Income shifting can take a myriad of different forms. For example, individuals may be able to shift income from personal to business tax bases through self-employment or incorporation. They may also have some facility to smooth reported income intertemporally by, for example, retaining earnings within a firm. Alternatively, income-shifting motives may also lead to some income remaining unreported (an evasion margin). In this paper, we focus on a particularly salient but somewhat underexplored channel of income shifting: *intra-family* income shifting. Many personal income tax systems are assessed at the individual level but allow households to influence how income is attributed across family members through business income allocation, trusts, beneficiary designations, and other legal arrangements. The family is therefore a natural candidate "syndicate" in which shifting occurs.

It is clear that income-shifting behaviors can engender very misleading interpretations of bunching elasticities. Indeed, even if taxpayers do not exhibit a fundamental earnings response to changes in tax rates, there could be substantial observed bunching simply by virtue of taxpayers adjusting their reported income through some income-shifting mechanism. A self-employed individual might, for example, retain just enough income within their firm to make their personal taxable income in the current period exactly coincide with a low tax

bracket. Alternatively, a high earner may shift their investment income to their lower-earning spouse to equalize marginal tax rates within the family and thereby lower the family's total tax burden.

Naturally, it would be worthwhile to determine what share of observed bunching reflects fundamental behavioral responses, that is, changes in economic activity, rather than income shifting. However, obtaining such a decomposition is a challenging task, since it requires an understanding of – as well as data pertaining to – a wide range of candidate income-shifting behaviors. Several papers in the literature have made progress in developing a more robust characterization of bunching elasticities, especially in terms of decomposing their underlying channels. For example, Le Maire and Schjerning (2013) extend Saez's bunching framework to a dynamic environment with the potential for intertemporal income smoothing through retained earnings and find that for the Danish self-employed, over half of observed bunching is explained by intertemporal income-shifting responses. Chetty et al. (2011) and Kleven and Waseem (2013) clarify the interpretation of bunching elasticities in the presence of adjustment costs, distinguishing between structural and observed elasticities.

Our main contribution in this paper is a falsification approach based on aggregation within what we call a *shiftable syndicate*. A shiftable syndicate is any collection of components across which reported income can be transferred while satisfying an aggregate accounting identity. In the case of intra-family income shifting, the components are family members and the accounting identity is that total family income equals the sum of members' reported incomes. For the case of intertemporal income shifting, the "syndicate members" are different time periods. We show that, under relatively mild conditions, *fundamental* bunching (that is, bunching in the absence of shifting possibilities) at each component's kink necessarily implies bunching of syndicate-level aggregate income ("multi-bunching"). We provide two theoretical existence results to validate this intuition, one with costless shifting and another with convex shifting costs. The reverse implication, which can be tested empirically, is that the absence of bunching at the syndicate level must also imply the absence of fundamental

bunching at the individual level.

We then take this implication to the data. We use Australian administrative tax microdata linked into family units. The Australian Taxation Office Longitudinal Information Files (ALife) provide a large, high-quality panel of personal income tax records (Abhayaratna et al., 2022). The recently developed ALife family file links individuals into spouse–child groupings in each year, creating an observed family unit suitable for testing the shiftable-syndicate prediction (Fischer et al., 2025). Australia is a useful laboratory because its personal income tax schedule features salient kink points in the middle of the distribution and prior work documents strong individual-level bunching (Johnson et al., 2024). Most importantly, the absence of transfer taxes and the widespread use of family trusts make the intra-family shifting channel particularly salient.

Empirically, rather than estimate a local bunching elasticity, we ask a complementary question: is an income distribution *globally* well approximated by a smooth density? Large administrative samples make visual diagnostics unreliable: with enough observations, tiny local deviations become visually striking. We therefore implement globally valid goodness-of-fit procedures that compare the empirical distribution to a nonparametrically fitted smooth benchmark. We consider both Cramér–von Mises (CvM) and Kolmogorov–Smirnov (KS) discrepancy measures, and we calibrate inference using (i) holdout sample splitting and (ii) a smooth bootstrap with re-fitting. This approach is conceptually related to density-based manipulation tests (McCrary, 2008; Cattaneo et al., 2020), but targeted to the global shape of the distribution rather than a discontinuity at a single cutoff.

Our headline finding is that individual taxable income displays sharp spikes at statutory kink points, but the corresponding family aggregates are remarkably smooth. The global smoothness null is strongly rejected for individuals in every year, using both CvM and KS and under both calibration methods. In contrast, the same procedures frequently fail to reject smoothness for family-level aggregates. This pattern is difficult to reconcile with widespread fundamental labor-supply responses at the kinks: if true resources adjusted to

locate at the kink, aggregation should preserve some trace of the excess mass. Instead, the evidence is consistent with the vast majority of individual bunching reflecting intra-family income shifting.

The remainder of the paper is organized as follows. Section 2 develops the shiftable syndicate framework and proves the aggregation implications. Section 3 describes the ALife individual and family data. Section 4 presents our global smoothness tests and calibration procedures. Section 5 summarizes the empirical evidence, and Section 6 concludes.

2 A shiftable syndicate framework

2.1 From individual bunching to syndicate bunching

A standard bunching model begins with an agent of type θ choosing effort e to produce income $y = f(e, \theta)$, facing a nonlinear tax schedule $T(\cdot)$ with a kink at y_k. Under regularity conditions, the optimal report y may exhibit excess mass at y_k. The key complication for interpretation is that the *reported* y may differ from the *true* resource $f(e, \theta)$ because agents can re-label income via accounting, transfers, or timing.

We therefore consider a collection of components $t = 1, \ldots, T$ that jointly satisfy an aggregate accounting identity. The intended empirical mapping is: t indexes family members (spouse, children) or time periods; x_t is the taxable income attributed to component t; and the aggregate resource is the sum of true incomes $\sum_t f_t(e_t, \theta_t)$. Shifting allows the allocation of the aggregate resource across components' tax bases.

Definition 1 (Shiftable syndicate)**.** *A* shiftable syndicate *consists of T components indexed by $t = 1, \ldots, T$. Component t produces true income $f_t(e_t, \theta_t)$ by choosing effort $e_t \geq 0$, and faces a tax schedule $T_t(\cdot)$ applied to its reported income $x_t \geq 0$. A* shifting arrangement *is a*

vector $s = (s_1, \ldots, s_T)$ *such that*

$$x_t = f_t(e_t, \theta_t) + s_t, \qquad \sum_{t=1}^{T} s_t = 0,$$

possibly subject to a convex shifting cost $S(s)$*.*

Definition 2 (Fundamental bunching)**.** *A type-level notion of* fundamental *bunching is defined relative to a counterfactual in which shifting is impossible. Fix a component* t *and a kink point* y_{t,j_t}*. Consider the* no-shifting *problem with* $s_t \equiv 0$*, so that reported income equals true income* ($x_t = f_t(e_t, \theta_t)$)*. We say that a type* θ_t *exhibits* fundamental bunching *at* y_{t,j_t} *if the global optimum of the no-shifting problem satisfies* $x_t^0(\theta_t) = y_{t,j_t}$*. Intuitively, absent any ability to reallocate resources across components, the kink itself pins down the optimum for that type.*

Under this definition, shifting re-allocates reported income across components without changing total resources. Intra-family shifting fits naturally: a family trust or beneficiary designation can change which individual's return records a given amount of income, while the family's total income is preserved up to compliance frictions. Intertemporal shifting is the same object with t indexing years. Now, the core falsification idea that we develop is the following: if individuals exhibit real responses, then we should observe bunching in both individual taxable income and syndicate-level aggregate taxable income. That is, we should observe nontrivial bunching at the syndicate level, where all syndicate members locate at their individual kink points simultaneously. It is worth noting that this would occur even if income shifting were completely impossible. As the capacity to shift income increases, the degree of observed bunching should also rise, since this mitigates the need for all members of the syndicate to exhibit fundamental bunching.[1] If, instead, individual-level bunching primarily reflects redistribution of reported income within the syndicate rather than real

[1] In the extreme case that the syndicate has sufficient buffer-stock income to redistribute and shifting is costless, we would see all but one individual in the syndicate locate at a kink point, requiring fundamental bunching only from the final remaining member of the syndicate.

responses, aggregation should largely eliminate the spikes.

2.2 Existence of multi-bunching equilibria

We now formalize the above intuition in two results: one with free shifting (no shifting cost) and one with costly shifting. The results are written for a general multi-component environment.

Proposition 1 (Free shifting)**.** *Let $T \in \mathbb{N}$. For each $t = 1, \ldots, T$, let $T_t : \mathbb{R}_+ \to \mathbb{R}_+$ be a convex piecewise linear tax function with a kink at $y_{t,j_t} > 0$, where the marginal tax rates to the left and right of the kink, respectively, satisfy $0 \leq \tau^-_{t,j_t} < \tau^+_{t,j_t} < 1$. Denote after-tax income by $h_t(x) \coloneqq x - T_t(x)$ (so that h_t is concave), and let the utility function over after-tax income $u : \mathbb{R}_+ \to \mathbb{R}$ be C^1, strictly increasing and strictly concave. Suppose also that the production function $f_t(\cdot, \theta_t)$ is C^1, strictly increasing and concave in $e \geq 0$, and that the effort cost $c(\cdot; \theta_t)$ is C^1, strictly increasing and strictly convex in $e \geq 0$. For $\theta = (\theta_1, \ldots, \theta_T)$, consider*

$$\max_{x,e} \sum_{t=1}^{T} \big[u(h_t(x_t)) - c(e_t; \theta_t)\big] \quad s.t. \quad \sum_{t=1}^{T} x_t = \sum_{t=1}^{T} f_t(e_t, \theta_t), \quad x_t, e_t \geq 0.$$

Define $c^*_t \coloneqq y_{t,j_t} - T_t(y_{t,j_t})$, $U_t \coloneqq u'(c^*_t)$, and the admissible multiplier interval at the kink

$$\Lambda_t \coloneqq \big[\, U_t(1 - \tau^+_{t,j_t}) \,,\, U_t(1 - \tau^-_{t,j_t}) \,\big] \subset (0, \infty).$$

For $\lambda > 0$, let $\psi_t(\lambda; \theta_t)$ be the unique solution to $c_e(\psi_t; \theta_t) = \lambda\, f_{t,e}(\psi_t, \theta_t)$, and set

$$Z(\lambda; \theta) \coloneqq \sum_{t=1}^{T} f_t(\psi_t(\lambda; \theta_t), \theta_t).$$

If there exist $\bar{\theta}$ and $\alpha < \beta$ with $[\alpha, \beta] \subset \bigcap_{t=1}^{T} \Lambda_t$ such that

$$Z(\alpha; \bar{\theta}) \;<\; \sum_{t=1}^{T} y_{t,j_t} \;<\; Z(\beta; \bar{\theta}),$$

then there is a nonempty open neighborhood U of $\bar{\theta}$ such that for every $\theta \in U$, the solution of the optimization problem is such that $x_t^* = y_{t,j_t}$ for all t.

Proof. By continuity of ψ_t and f_t in (λ, θ_t), the strict inequalities persist on some open neighborhood U of $\bar{\theta}$. Fix $\theta \in U$. The map $\lambda \mapsto Z(\lambda; \theta)$ is continuous and strictly increasing, hence there exists a unique $\lambda(\theta) \in (\alpha, \beta)$ with $Z(\lambda(\theta); \theta) = \sum_t y_{t,j_t}$. Let $e_t^* \coloneqq \psi_t(\lambda(\theta); \theta_t)$ and $x_t^* \coloneqq y_{t,j_t}$. Then feasibility holds by construction.

Consider the Lagrangian

$$\mathcal{L}(e, x; \lambda) = \sum_{t=1}^{T} \big(u(h_t(x_t)) - \lambda x_t\big) + \sum_{t=1}^{T} \big(\lambda f_t(e_t, \theta_t) - c(e_t; \theta_t)\big),$$

with $\lambda = \lambda(\theta)$. For each t, h_t is concave and piecewise linear with left/right slopes $1 - \tau_{t,j_t}^{+}$ and $1 - \tau_{t,j_t}^{-}$ at y_{t,j_t}, so

$$\partial h_t(y_{t,j_t}) = [\, 1 - \tau_{t,j_t}^{+},\, 1 - \tau_{t,j_t}^{-} \,].$$

Since u is C^1, the chain rule for subdifferentials yields

$$\partial(u \circ h_t)(y_{t,j_t}) = u'(h_t(y_{t,j_t}))\, \partial h_t(y_{t,j_t}) = U_t\, \partial h_t(y_{t,j_t}) = \Lambda_t.$$

Thus the x–stationarity condition holds at $x_t^* = y_{t,j_t}$ if $\lambda \in \Lambda_t$, which is satisfied because $\lambda(\theta) \in (\alpha, \beta) \subset \Lambda_t$. The e–stationarity condition holds by definition of ψ_t. Our convexity assumptions then imply that (e^*, x^*) is globally optimal and that this solution is unique. □

Proposition 2 (Costly shifting)**.** *Maintain the objects u, T_t, f_t, c from Proposition 1. Introduce shifts $s_t \in \mathbb{R}$ with $\sum_{t=1}^{T} s_t = 0$ and separable shifting cost $S(s) = \sum_{t=1}^{T} \phi_t(s_t)$, where*

each $\phi_t : \mathbb{R} \to \mathbb{R}_+$ is C^1, strictly convex, and ϕ_t' is strictly increasing and onto $\mathbb{R}$. Consider

$$\max_{x,e,s} \sum_{t=1}^{T} \big[u(h_t(x_t)) - c(e_t;\theta_t)\big] \; - \; \sum_{t=1}^{T} \phi_t(s_t) \quad s.t. \quad x_t = f_t(e_t,\theta_t) + s_t, \;\; \sum_{t=1}^{T} s_t = 0.$$

Fix kinks y_{t,j_t} and define Λ_t as above. For each t, define the required shift at the kink

$$F_t(\lambda_t;\theta_t) \coloneqq y_{t,j_t} - f_t(\psi_t(\lambda_t;\theta_t),\theta_t).$$

For $\lambda = (\lambda_1,\dots,\lambda_T)$, let $\mu(\lambda)$ be the unique real number solving the equilibrium condition

$$\sum_{t=1}^{T} (\phi_t')^{-1}(\lambda_t + \mu) = 0.$$

Define $G : \prod_{t=1}^{T} \Lambda_t \to \mathbb{R}^T$ by

$$G_t(\lambda;\theta) \coloneqq (\phi_t')^{-1}(\lambda_t + \mu(\lambda)) - F_t(\lambda_t;\theta_t), \qquad t = 1,\dots,T.$$

Suppose there exist $\bar\theta$ and intervals $[\alpha_t,\beta_t] \subset \Lambda_t$ such that on the box $B = \prod_{t=1}^{T}[\alpha_t,\beta_t]$ the face-sign (Poincaré–Miranda) conditions hold: for each t,

$$G_t(\lambda;\bar\theta) \geq 0 \;\; \forall\, \lambda \in B \text{ with } \lambda_t = \alpha_t, \qquad G_t(\lambda;\bar\theta) \leq 0 \;\; \forall\, \lambda \in B \text{ with } \lambda_t = \beta_t,$$

and the inequalities are strict at some point on each respective face. Then there is a nonempty open neighborhood U of $\bar\theta$ such that for every $\theta \in U$ the problem admits a global optimum with $x_t^* = y_{t,j_t}$ for all t.

Proof. The objective is concave and the constraints are affine; hence KKT conditions are sufficient. Let λ_t be the multiplier on $x_t = f_t(e_t,\theta_t) + s_t$ and let μ be the multiplier on $\sum_t s_t = 0$. The KKT system is: (i) x–stationarity: $0 \in \partial_{x_t}\big(u(h_t(x_t)) - \lambda_t x_t\big)$, which at $x_t = y_{t,j_t}$ is equivalent to $\lambda_t \in \Lambda_t$; (ii) e–stationarity: $c_e(e_t;\theta_t) = \lambda_t f_{t,e}(e_t,\theta_t)$, equivalent to

$e_t = \psi_t(\lambda_t; \theta_t)$; (iii) s–stationarity: $\phi_t'(s_t) = \lambda_t + \mu$, equivalent to $s_t = (\phi_t')^{-1}(\lambda_t + \mu)$; (iv) feasibility: $x_t = f_t(e_t, \theta_t) + s_t$ and $\sum_t s_t = 0$.

Impose $x_t = y_{t,j_t}$. Then feasibility implies $s_t = F_t(\lambda_t; \theta_t)$, while shift stationarity implies $s_t = (\phi_t')^{-1}(\lambda_t + \mu)$. Eliminating s_t yields $G(\lambda; \theta) = 0$ and $\mu = \mu(\lambda)$.

By the Poincaré–Miranda theorem, the face-sign conditions imply the existence of λ^* in the interior of B with $G(\lambda^*; \bar{\theta}) = 0$. Set $\mu^* = \mu(\lambda^*)$, define

$$e_t^* = \psi_t(\lambda_t^*; \bar{\theta}_t), \qquad s_t^* = (\phi_t')^{-1}(\lambda_t^* + \mu^*), \qquad x_t^* = y_{t,j_t}.$$

Then feasibility and stationarity hold, so by concavity, (x^*, e^*, s^*) is the global optimum. Continuity of G in θ implies that the face-sign pattern persists on a neighborhood U of $\bar{\theta}$, yielding the desired result. □

2.3 Interpretation of theoretical results

Propositions 1 and 2 express the same economic force in two institutional regimes. With free shifting, a single multiplier λ governs the allocation of total resources across components. Each kink admits an *interval* of supporting multipliers Λ_t, because the objective is locally linear in x_t at the kink. The intersection $\cap_t \Lambda_t$ is therefore the set of aggregate shadow values that can support *every* component sitting at its kink. The bracketing condition on $Z(\lambda; \theta)$ requires that total produced income move continuously from below to above the sum of kink points as λ varies within that intersection. When this holds, the intermediate value theorem yields a multiplier that simultaneously (i) keeps all components at their kink points and (ii) satisfies the aggregate accounting identity.

With costly shifting, each component has its own local shadow value λ_t, and the shifting technology links these local shadow values through what is analogous to a market-clearing condition. The mapping $G(\lambda; \theta)$ summarizes the difference between the shifting *supply* induced by marginal shifting costs and the shifting *demand* required to place each component

exactly at its kink. The Poincaré–Miranda face-sign conditions are a multidimensional bracketing condition: on each face of the box, shifting supply exceeds demand on one side and falls short on the other, forcing an interior zero. Economically, these conditions state that for each component, there exist shadow values that generate too little shifting and too much shifting, and that these regions overlap across components. When they do, there exists a joint configuration of shadow values that clears the aggregate resource constraint while supporting multi-bunching.

These results connect *fundamental* bunching in the sense of Definition 2 to a sharp empirical implication. If, for each component, there are types who would optimally locate at the kink even when shifting is impossible and there is sufficient responsiveness of effort to changes in the tax schedule, then our propositions show that (under the stated overlap and bracketing conditions) there exists a positive measure of equilibria in which *all* components can be placed at their kink points simultaneously. Because the syndicate's aggregate resource constraint then forces the total to equal the sum of kink locations, the syndicate aggregate inherits excess mass ("multi-bunching"). Conversely, if observed individual spikes are generated solely through *ex-post* reallocation of a fixed total across components (pure within-syndicate relabeling), then aggregation to the syndicate total should largely eliminate the spikes.

3 Data

3.1 Administrative tax data and family linkage

Our analysis uses Australian administrative income tax data collected by the Australian Taxation Office (ATO). The ALife dataset is a 10% longitudinal sample of individuals' personal income tax returns, designed for policy and research use (Abhayaratna et al., 2022). We use the ALife family file, which links individuals to spouses and children (as observed in administrative records) to form family units (Fischer et al., 2025). The family file is

particularly useful for studying joint financial decisions and intergenerational outcomes, but in this context, it also enables a direct test of intra-family income shifting.

We work at the annual tax-return level. Let y_{it} denote individual i's taxable income in year t (in nominal dollars), and let $g(i,t)$ denote the family (syndicate) to which individual i belongs in year t according to the family file. We then perform the following constructions:

- **Family group:** We construct, for each individual, a family group consisting of that individual, their spouse, and their children. To avoid spurious aggregation over excessively large family groups, we restrict families to have at most two generations.
- **Family income:** the family's total taxable income is the sum of taxable incomes across the individuals in the family group.
- **Average income:** to ensure that incomes are on the same scale for different family groups, we largely consider total family income divided by the number of family members (averaging makes no difference to the theoretical results discussed in the previous section).

Sample definition. A practical challenge is that administrative family linkages are incomplete: some individuals appear as having spouse income fields without a corresponding spouse identifier in the relationships file. To mitigate unobserved shifting, we adopt a conservative sample rule: we exclude "singletons" (individuals with no recorded spouse or child links in the family file) and restrict attention to family units for which the linkage is internally consistent. We choose to focus on families in the range where kink points are relevant (average income between \$0 and \$200,000, since the top bracket is less than \$200,000 in all years).

4 Empirical strategy: testing global smoothness

4.1 Global non-bunching as a goodness-of-fit null

Fix a year and an income concept (individual taxable income, couple total income, family total income, or per-individual average income within a family). Let $X_1, \ldots, X_n$ be the sample of incomes. We interpret "global non-bunching" as the hypothesis that X is drawn from a distribution with a smooth density that is well captured by a nonparametric kernel density estimate (KDE). Under this null, a smooth fitted CDF should approximate the empirical CDF well. Under alternatives with bunching, the empirical CDF exhibits localized steepening at kink points, leading to large discrepancies. The rationale for this global approach is that there may be many candidate composite bunching points with aggregation, making such a test more informative about the global characteristics of the income distribution.

Let $\hat{F}$ denote the fitted CDF implied by a KDE. For a sample $x_1, \ldots, x_m$, let F_m denote the empirical CDF. We consider two discrepancy measures:

$$T^{\mathrm{CvM}} = \int \left(F_m(x) - \hat{F}(x)\right)^2 d\hat{F}(x),$$
$$T^{\mathrm{KS}} = \sup_x \left|F_m(x) - \hat{F}(x)\right|.$$

Both statistics are sensitive to departures from the fitted smooth distribution; the Kolmogorov–Smirnov statistic is particularly sensitive to localized spikes, while the Cramér–von Mises statistic averages squared deviations and can be more robust to small local noise.

A naive application (fit $\hat{F}$ on the full sample and then test on the same sample) produces misleading results, because $\hat{F}$ is optimized on the same data being evaluated. We therefore propose two ways to restore validity: holdout sample-splitting and smooth bootstrap re-fitting.

4.2 Method 1: holdout (sample-splitting) tests

Randomly partition the sample into a training set A of size n_A and a test set B of size n_B. Fit a KDE on A to obtain $\hat{f}_A$ and its implied CDF $\hat{F}_A$. Then compute a one-sample goodness-of-fit statistic on B against the fixed reference $\hat{F}_A$:

$$T_{B|A}^{\mathrm{CvM}} = \int \left(F_B(x) - \hat{F}_A(x)\right)^2 d\hat{F}_A(x),$$
$$T_{B|A}^{\mathrm{KS}} = \sup_x \left|F_B(x) - \hat{F}_A(x)\right|.$$

Conditional on the training sample, this is a standard one-sample goodness-of-fit problem. As $n_A \to \infty$, $\hat{F}_A$ converges to the underlying smooth target under the null, and the resulting test has asymptotically correct size.

4.3 Method 2: smooth bootstrap with re-fitting

The holdout approach is asymptotically valid but can be noisy in finite samples because $\hat{F}_A$ is itself estimated. We therefore also implement a *smooth bootstrap* protocol aimed at finite-sample calibration. Let $\hat{f}_0$ denote the KDE fitted to the full sample and let $\hat{F}_0$ be its implied CDF. We compute the observed discrepancies T_{obs}^{CvM} and T_{obs}^{KS} between the empirical CDF of the original sample, F_n, and the fitted CDF $\hat{F}_0$. To approximate the null distribution, for $b = 1, \ldots, B$:

1. Draw a bootstrap sample $X_1^{*(b)}, \ldots, X_n^{*(b)}$ from $\hat{f}_0$ (equivalently, resample observed points and add kernel noise).

2. Re-fit the KDE on the bootstrap sample using the same bandwidth selection rule to obtain $\hat{f}_0^{*(b)}$ and $\hat{F}_0^{*(b)}$.

3. Compute T_b^{CvM} and T_b^{KS} by comparing the bootstrap empirical CDF $F_n^{*(b)}$ to $\hat{F}_0^{*(b)}$.

The bootstrap p-value is the exceedance probability

$$\hat{p} = \frac{1}{B} \sum_{b=1}^{B} \mathbb{1}\{T_b \geq T_{\text{obs}}\}.$$

4.4 Implementation details

We implement KDE fits using a Gaussian kernel with automatic bandwidth selection based on the plug-in rule developed by Sheather and Jones (1991). In practice, taxable income is recorded in dollars and often has mass points due to rounding and reporting conventions. KS and CvM statistics are well-defined with ties, but their classical null distributions assume continuity. We therefore adopt the safeguard of adding mean-zero jitter that is small relative to the reporting unit (e.g., Uniform$(-\epsilon, \epsilon)$ with $\epsilon < 10$) before computing the test statistics.

5 Results

This section summarizes the empirical patterns that we observe in the data. To keep the main text concise, we highlight an exemplar tax year in figures, while reporting year-by-year goodness-of-fit outcomes (2009–2022) in tables; the full set of diagnostic figures by year is relegated to the Appendix. Here, we illustrate the results for a given exemplar year of 2021.

5.1 Bunching in individual taxable income at kink points

Consistent with prior evidence for Australia (Johnson et al., 2024), individual taxable income distributions show visually striking spikes at bracket thresholds (marked on the figure with vertical dotted lines). Under the standard bunching interpretation, these patterns would suggest positive ETI at multiple kink points. Figure 1 shows the histogram for individual taxable income with a KDE overlay in the exemplar year 2021.

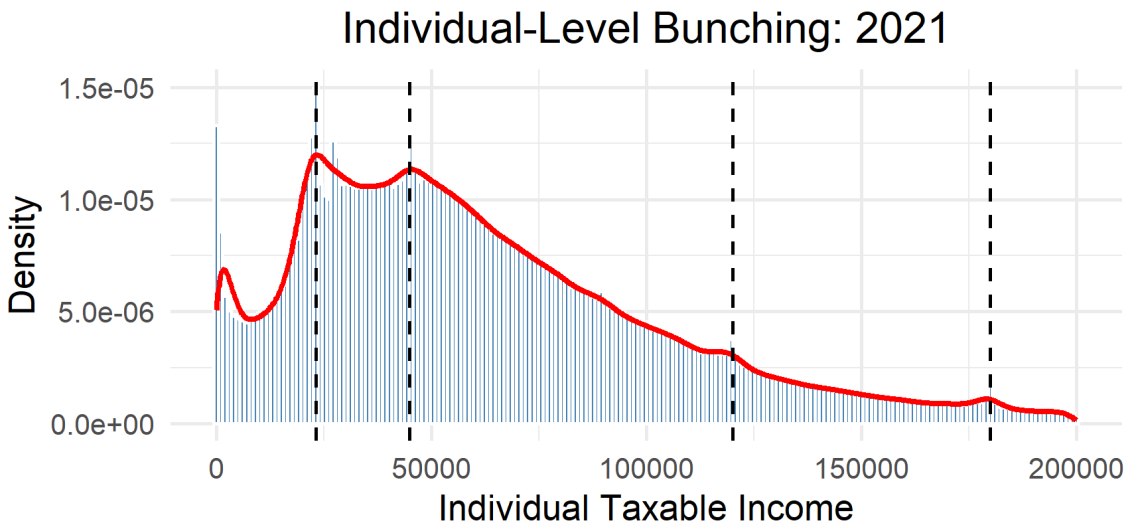


Figure 1: Histogram and KDE overlay plot for individual taxable income, tax year 2021.

5.2 Aggregation to families eliminates visible bunching

When we aggregate taxable income to the family level, the spikes are no longer present. Average per-individual family income exhibits smooth distributions that, in visual diagnostics, are well approximated by a globally smooth density. This is consistent with the shiftable syndicate logic: if much of the individual-level bunching is driven by intra-family allocation of income across returns, then aggregation should undo the relabeling and restore smoothness. Figure 2 reports the corresponding KDE overlay diagnostic for average per-individual family income in the exemplar year 2021 (excluding singletons, as noted earlier).

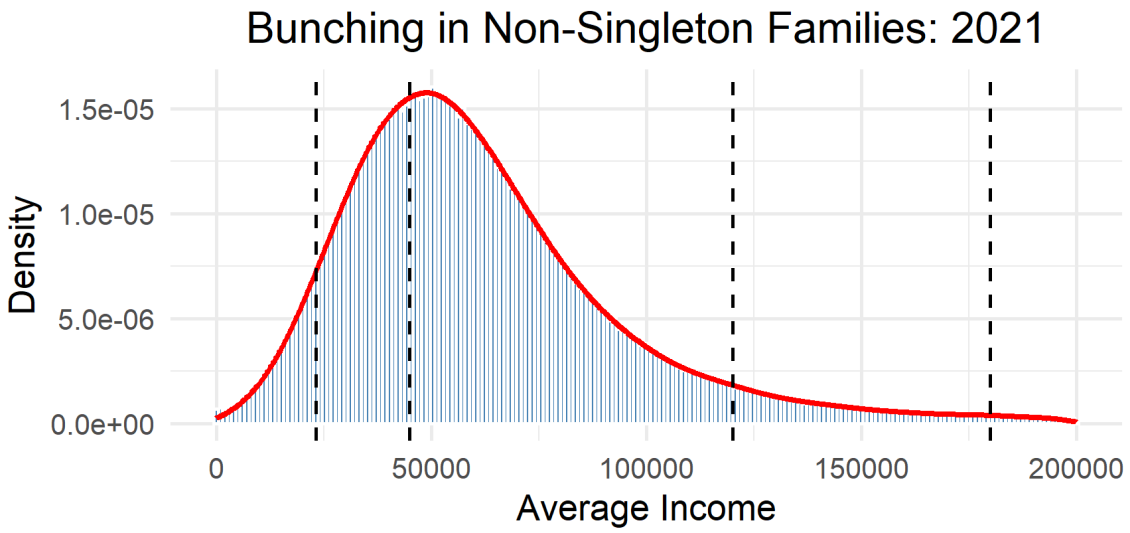


Figure 2: Histogram and KDE overlay plot for average per-individual family income, tax year 2021.

5.3 Global smoothness tests

Our goodness-of-fit tests confirm the visual intuition. For individuals, both the holdout and re-fit bootstrap procedures strongly reject global smoothness in every year, using both CvM and KS statistics. For families, by contrast, we fail to reject smoothness in several years (particularly using the bootstrap KS test), and the distributional discrepancies are an order of magnitude smaller. It is worth noting that the KS and CvM tests are very powerful for sample sizes as large as those in this paper, in that even very small distributional deviations can lead to rejection of the null. To this end, even the fact that we are unable to reject the null in several years should be construed as compelling evidence of distributions that are close to being globally smooth.

Tables 1 and 2 report the CvM and KS statistics and associated p-values by tax year (2009–2022), for individuals and for family aggregates. Columns three through six implement method one above, and columns seven through ten implement method two.

Table 1: CvM goodness-of-fit tests by year. Superscripts "hold" and "boot" refer to tests conducted using the holdout and bootstrap methods, respectively. Subscripts I and F denote individual- and family-level tests, respectively.

Year	N	$\text{Stat}_I^{\text{hold}}$	p_I^{hold}	$\text{Stat}_F^{\text{hold}}$	p_F^{hold}	$\text{Stat}_I^{\text{boot}}$	p_I^{boot}	$\text{Stat}_F^{\text{boot}}$	p_F^{boot}
2009	488,043	4.083	$< 10^{-4}$	1.426	0.0003	8.226	$< 10^{-4}$	0.869	$< 10^{-4}$
2010	487,589	3.711	$< 10^{-4}$	2.137	$< 10^{-4}$	8.733	$< 10^{-4}$	0.869	$< 10^{-4}$
2011	488,736	4.358	$< 10^{-4}$	0.348	0.100	8.129	$< 10^{-4}$	0.787	$< 10^{-4}$
2012	481,867	3.434	$< 10^{-4}$	0.460	0.050	8.132	$< 10^{-4}$	0.619	$< 10^{-4}$
2013	473,173	3.610	$< 10^{-4}$	0.390	0.077	7.715	$< 10^{-4}$	0.546	0.007
2014	473,971	4.536	$< 10^{-4}$	1.313	0.0005	8.449	$< 10^{-4}$	0.572	0.003
2015	474,515	4.964	$< 10^{-4}$	1.166	0.001	9.134	$< 10^{-4}$	0.533	0.013
2016	473,484	5.056	$< 10^{-4}$	1.229	0.0007	9.562	$< 10^{-4}$	0.540	$< 10^{-4}$
2017	471,997	4.399	$< 10^{-4}$	0.568	0.027	9.574	$< 10^{-4}$	0.540	$< 10^{-4}$
2018	466,952	4.712	$< 10^{-4}$	0.212	0.245	9.610	$< 10^{-4}$	0.525	$< 10^{-4}$
2019	460,543	3.307	$< 10^{-4}$	2.057	$< 10^{-4}$	9.063	$< 10^{-4}$	0.483	$< 10^{-4}$
2020	450,380	3.368	$< 10^{-4}$	0.423	0.063	7.241	$< 10^{-4}$	0.425	0.010
2021	438,295	2.229	$< 10^{-4}$	0.093	0.620	4.537	$< 10^{-4}$	0.309	0.623
2022	422,554	3.643	$< 10^{-4}$	0.374	0.085	6.944	$< 10^{-4}$	0.399	0.063

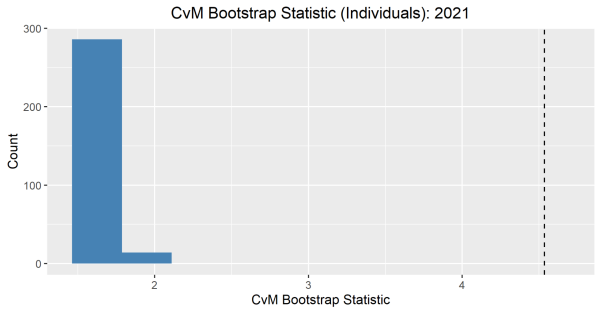
Table 2: KS goodness-of-fit tests by year. Superscripts "hold" and "boot" refer to tests conducted using the holdout and bootstrap methods, respectively. Subscripts I and F denote individual- and family-level tests, respectively.

Year	N	$\text{Stat}_I^{\text{hold}}$	p_I^{hold}	$\text{Stat}_F^{\text{hold}}$	p_F^{hold}	$\text{Stat}_I^{\text{boot}}$	p_I^{boot}	$\text{Stat}_F^{\text{boot}}$	p_F^{boot}
2009	488,043	0.0058	$< 10^{-4}$	0.0047	$< 10^{-4}$	0.0059	$< 10^{-4}$	0.0023	0.010
2010	487,589	0.0066	$< 10^{-4}$	0.0039	0.001	0.0064	$< 10^{-4}$	0.0023	0.003
2011	488,736	0.0062	$< 10^{-4}$	0.0069	$< 10^{-4}$	0.0061	$< 10^{-4}$	0.0023	$< 10^{-4}$
2012	481,867	0.0062	$< 10^{-4}$	0.0024	0.117	0.0065	$< 10^{-4}$	0.0021	0.067
2013	473,173	0.0067	$< 10^{-4}$	0.0027	0.059	0.0067	$< 10^{-4}$	0.0021	0.020
2014	473,971	0.0071	$< 10^{-4}$	0.0027	0.059	0.0070	$< 10^{-4}$	0.0020	0.130
2015	474,515	0.0078	$< 10^{-4}$	0.0027	0.068	0.0078	$< 10^{-4}$	0.0019	0.390
2016	473,484	0.0079	$< 10^{-4}$	0.0057	$< 10^{-4}$	0.0080	$< 10^{-4}$	0.0020	0.100
2017	471,997	0.0081	$< 10^{-4}$	0.0023	0.153	0.0080	$< 10^{-4}$	0.0020	0.017
2018	466,952	0.0080	$< 10^{-4}$	0.0017	0.484	0.0081	$< 10^{-4}$	0.0019	0.303
2019	460,543	0.0079	$< 10^{-4}$	0.0036	0.005	0.0078	$< 10^{-4}$	0.0018	0.310
2020	450,380	0.0074	$< 10^{-4}$	0.0015	0.696	0.0073	$< 10^{-4}$	0.0018	0.310
2021	438,295	0.0071	$< 10^{-4}$	0.0029	0.046	0.0070	$< 10^{-4}$	0.0016	0.447
2022	422,554	0.0071	$< 10^{-4}$	0.0028	0.077	0.0072	$< 10^{-4}$	0.0017	0.763

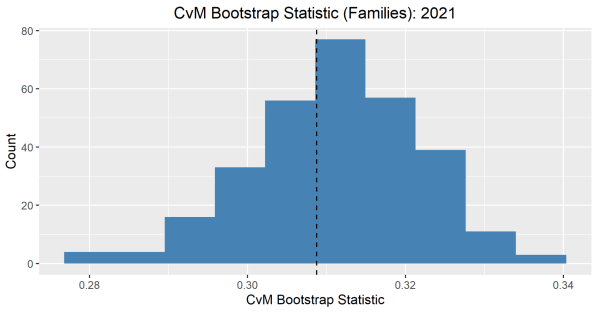
Table 3: Exemplar-year goodness-of-fit results (2021). The table reports p-values from (i) holdout sample-splitting and (ii) smooth bootstrap with re-fitting.

Test	N	p_I^{hold}	p_F^{hold}	p_I^{boot}	p_F^{boot}
CvM	438,295	$< 10^{-4}$	0.620	$< 10^{-4}$	0.623
KS	438,295	$< 10^{-4}$	0.046	$< 10^{-4}$	0.447

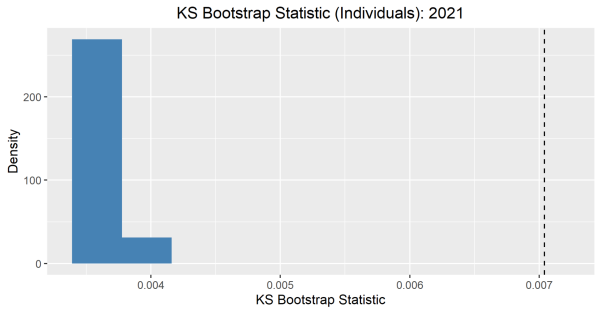
(a) CvM bootstrap (individuals)



(b) CvM bootstrap (families)



(c) KS bootstrap (individuals)



(d) KS bootstrap (families)

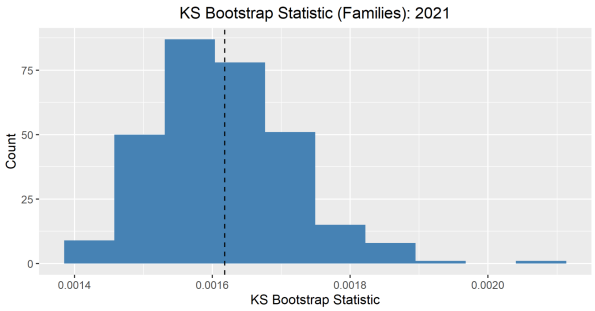


Figure 3: Smooth bootstrap with re-fitting in the exemplar year (2021). Each panel shows the bootstrap distribution of the test statistic under the fitted KDE null (re-fitting within each replicate); the dashed line marks the observed statistic.

5.4 Interpretation and limitations

The results highlight that ETI estimates based on individual taxable income may primarily reflect shifting opportunities rather than real effort responses in the Australian context. Naturally, we have only considered one type of income-shifting behavior here, but it is entirely possible that shifting may also occur in other directions (across firms, across tax bases, across years). However, intra-family shifting appears to explain much of the observed

bunching in individual taxable income for Australian taxpayers.

6 Conclusion

Bunching methods are a powerful tool to learn about taxpayers' responses to nonlinear incentives, but their interpretation can be fragile in the presence of income-shifting opportunities. This paper proposes a simple falsification approach based on aggregation within *shiftable syndicates*. Under general conditions, fundamental bunching at the component level should imply bunching in syndicate aggregates; when it does not, the component-level spikes are likely driven by within-syndicate reallocation.

Applying this logic to Australian tax data linked into family units, we find a sharp contrast: individual taxable income shows pronounced spikes at kink points, but total family income is largely smooth. Global goodness-of-fit tests using holdout sample-splitting and smooth bootstrap re-fitting, and using both CvM and KS discrepancy measures, corroborate this conclusion across years. The results suggest that intra-family income shifting may account for a substantial share of the individual-level bunching that would otherwise be interpreted as a real behavioral response.

For policy, the findings caution against interpreting individual bunching as evidence of large real ETI in contexts where income can be reassigned within families. A caveat of our results is that the predominance of intra-family income shifting is likely specific to the Australian context, driven by the substantial use of family trusts and the absence of transfer taxes. However, our shiftable syndicate perspective can be applied to other environments where relabeling is feasible, such as shifting between corporate and personal tax bases or intertemporal shifting. In each case, identifying the relevant syndicate and testing whether bunching persists under aggregation provides a useful diagnostic for distinguishing real responses from reallocation. Normatively, our results suggest that policymakers should focus on the effective shadow rates that apply at the syndicate level, not necessarily the statutory

rates that govern the individual tax schedule. There may be opportunities to raise rates overall at a low efficiency cost, but reforms that sharply increase the progressivity of the tax system at the individual level may be self-defeating through income-shifting channels.

References

Abhayaratna, T., A. Carter, and S. Johnson (2022): "The ATO Longitudinal Information Files (ALife): Individuals—A new dataset for public policy research," *Australian Economic Review*, 55, 541–557.

Cattaneo, M. D., M. Jansson, and X. Ma (2020): "Simple local polynomial density estimators," *Journal of the American Statistical Association*, 115, 1449–1455.

Chetty, R., J. N. Friedman, T. Olsen, and L. Pistaferri (2011): "Adjustment costs, firm responses, and micro vs. macro labor supply elasticities: Evidence from Danish tax records," *The Quarterly Journal of Economics*, 126, 749–804.

Feldstein, M. (1999): "Tax avoidance and the deadweight loss of the income tax," *Review of Economics and Statistics*, 81, 674–680.

Fischer, D., J. Holland, M. Chan, C. Polidano, H. Vu, and R. Wilkins (2025): "The ATO ALife Family File," *Australian Economic Review*.

Gordon, R. H. and J. Slemrod (1998): "Are "Real" Responses to Taxes Simply Income Shifting between Corporate and Personal Tax Bases?" NBER Working Paper 6576, National Bureau of Economic Research, Cambridge, MA.

Johnson, S., R. Breunig, M. Olivo-Villabrille, and A. Zaresani (2024): "Individuals' responsiveness to marginal tax rates: Evidence from bunching in the Australian personal income tax," *Labour Economics*, 87, 102461.

Kleven, H. J. and M. Waseem (2013): "Using notches to uncover optimization frictions and structural elasticities: Theory and evidence from Pakistan," *The Quarterly Journal of Economics*, 128, 669–723.

Kopczuk, W. (2025): "Shifting and Other Problems with Taxable Income Elasticity: Joel Slemrod's Contributions to What We Know About Taxing Ourselves," *National Tax Journal*, 78, 751–766.

Le Maire, D. and B. Schjerning (2013): "Tax bunching, income shifting and self-employment," *Journal of Public Economics*, 107, 1–18.

McCrary, J. (2008): "Manipulation of the running variable in the regression discontinuity design: A density test," *Journal of Econometrics*, 142, 698–714.

Saez, E. (2010): "Do taxpayers bunch at kink points?" *American Economic Journal: Economic Policy*, 2, 180–212.

Sheather, S. J. and M. C. Jones (1991): "A reliable data-based bandwidth selection method for kernel density estimation," *Journal of the Royal Statistical Society: Series B (Methodological)*, 53, 683–690.

A Additional details on the tests

This Appendix records implementation details and additional results that are omitted from the main text for brevity.

A.1 Practical computation of CvM and KS statistics

In practice, with a given CDF $\hat{F}$ and test sample $x_{(1)} \leq \cdots \leq x_{(m)}$, the one-sample CvM statistic can be computed as

$$T^{\mathrm{CvM}} = \frac{1}{12m} + \sum_{j=1}^{m} \left(\hat{F}(x_{(j)}) - \frac{2j-1}{2m} \right)^2,$$

and the KS statistic as

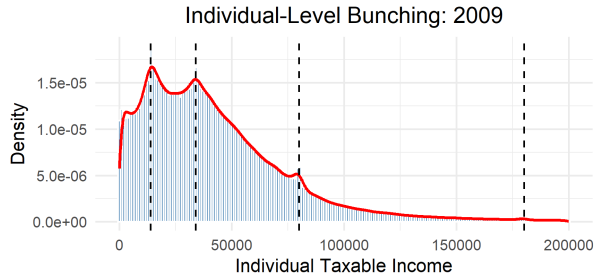
$$T^{\mathrm{KS}} = \max_{1 \leq j \leq m} \left\{ \left| \hat{F}(x_{(j)}) - \frac{j}{m} \right|, \left| \hat{F}(x_{(j)}) - \frac{j-1}{m} \right| \right\}.$$

For the holdout procedure, $\hat{F}$ is the training-set KDE CDF; for the bootstrap procedure, $\hat{F}$ is the re-fitted KDE CDF within each replicate.

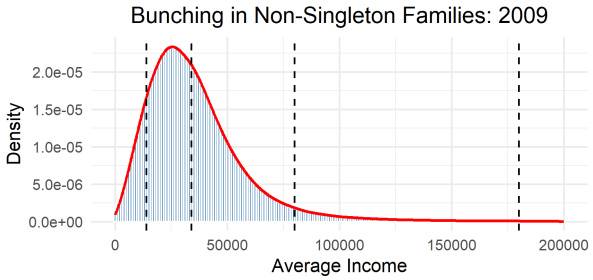
B Full time-series results

This Appendix collects the diagnostic figures referenced in the main text, reported year by year. Each page corresponds to one tax year.

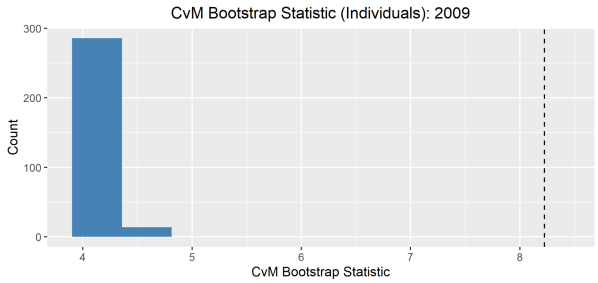
(a) Histogram and KDE overlay (individuals)



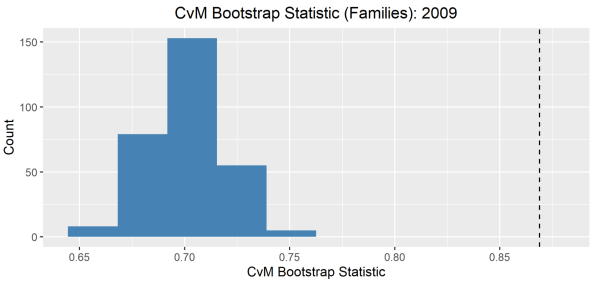
(b) Histogram and KDE overlay (families)



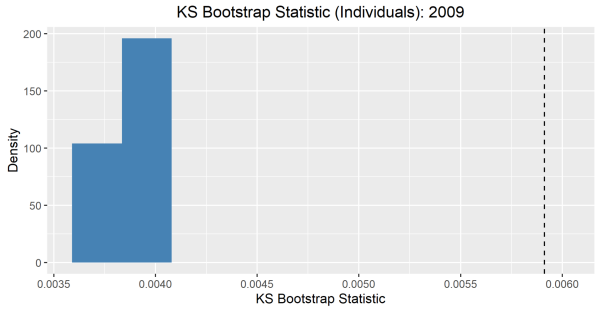
(c) CvM bootstrap (individuals)



(d) CvM bootstrap (families)



(e) KS bootstrap (individuals)



(f) KS bootstrap (families)

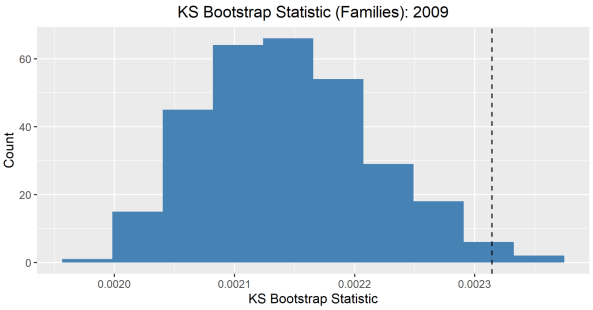
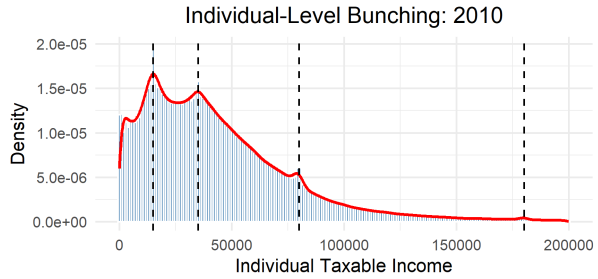
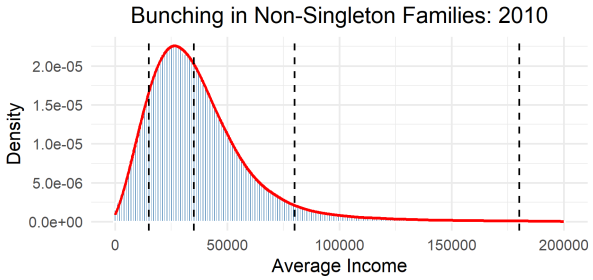


Figure 4: Year 2009: Histogram and KDE overlay with tax-threshold markers (top) and smooth-bootstrap goodness-of-fit diagnostics with re-fitting (bottom).

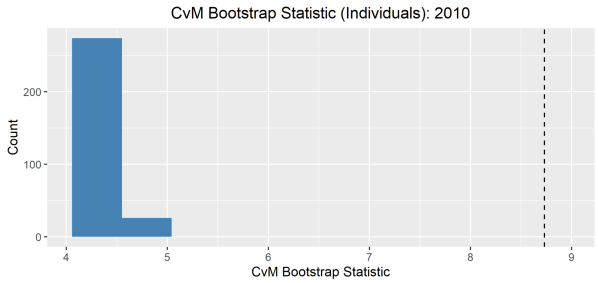
(a) Histogram and KDE overlay (individuals)



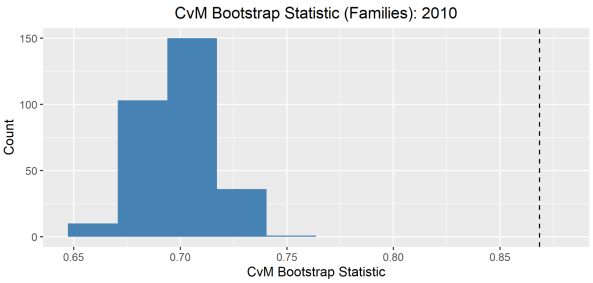
(b) Histogram and KDE overlay (families)



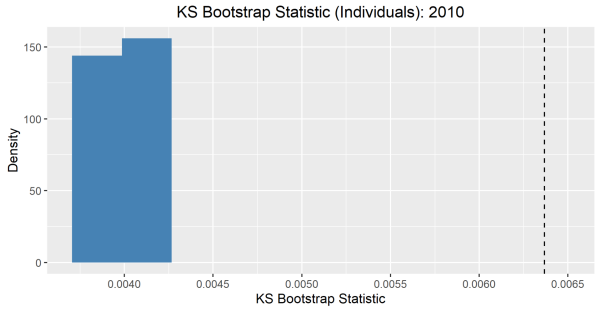
(c) CvM bootstrap (individuals)



(d) CvM bootstrap (families)



(e) KS bootstrap (individuals)



(f) KS bootstrap (families)

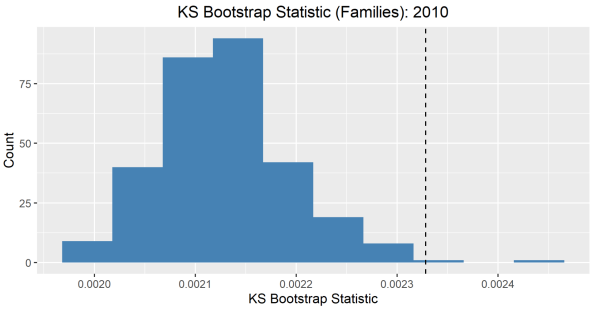
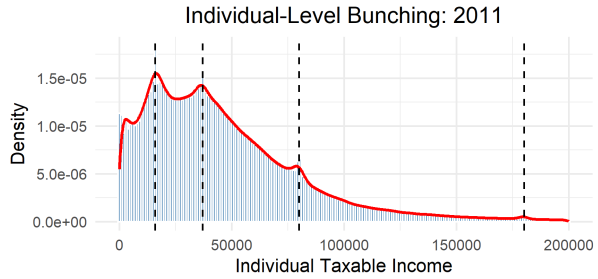
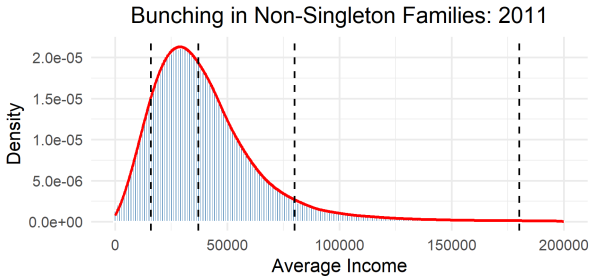


Figure 5: Year 2010: Histogram and KDE overlay with tax-threshold markers (top) and smooth-bootstrap goodness-of-fit diagnostics with re-fitting (bottom).

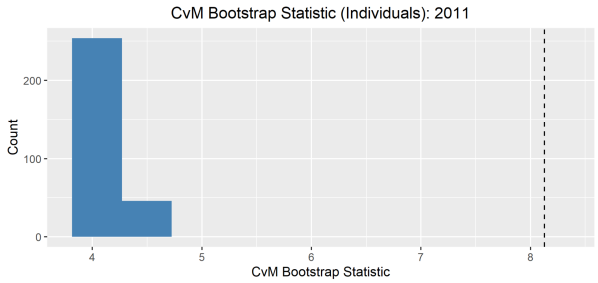
(a) Histogram and KDE overlay (individuals)



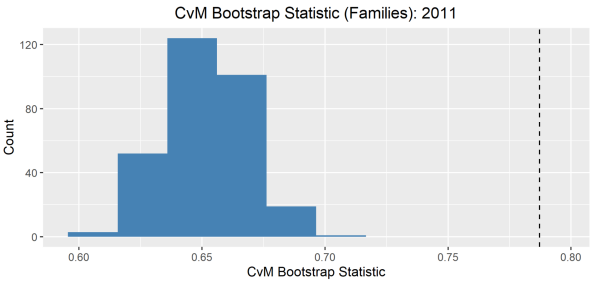
(b) Histogram and KDE overlay (families)



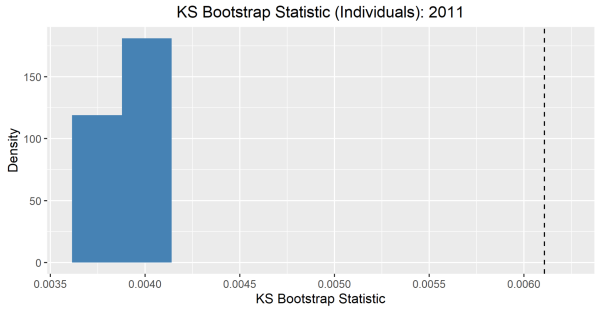
(c) CvM bootstrap (individuals)



(d) CvM bootstrap (families)



(e) KS bootstrap (individuals)



(f) KS bootstrap (families)

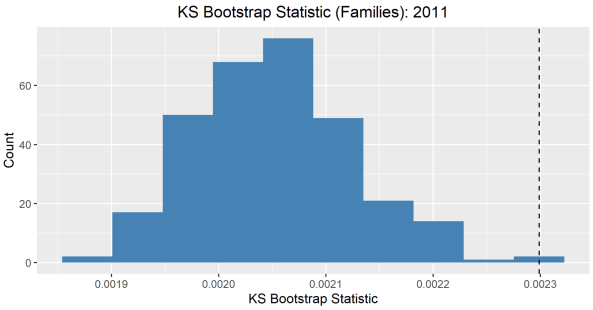
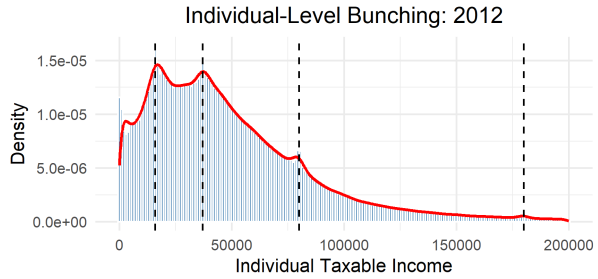
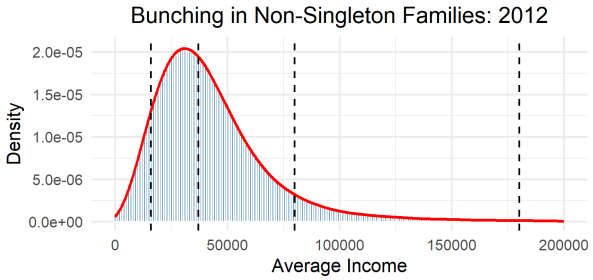


Figure 6: Year 2011: Histogram and KDE overlay with tax-threshold markers (top) and smooth-bootstrap goodness-of-fit diagnostics with re-fitting (bottom).

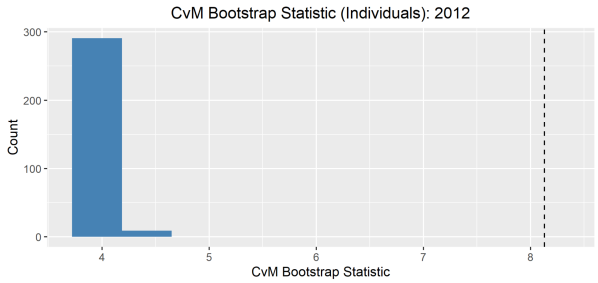
(a) Histogram and KDE overlay (individuals)



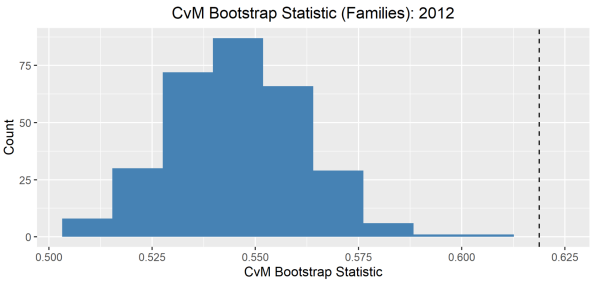
(b) Histogram and KDE overlay (families)



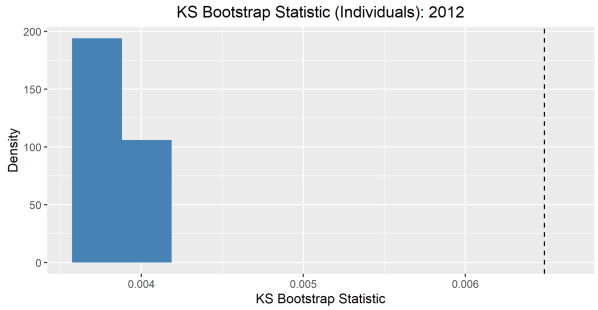
(c) CvM bootstrap (individuals)



(d) CvM bootstrap (families)



(e) KS bootstrap (individuals)



(f) KS bootstrap (families)

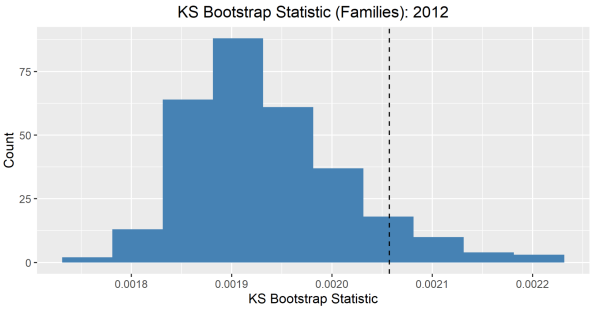
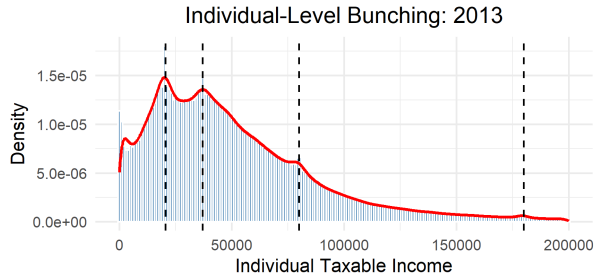
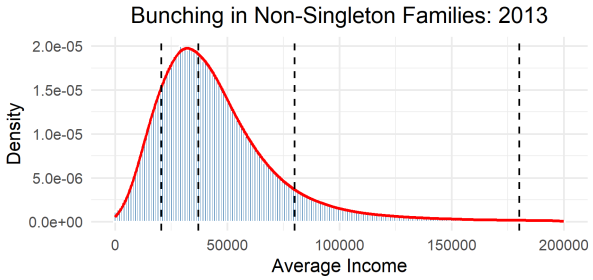


Figure 7: Year 2012: Histogram and KDE overlay with tax-threshold markers (top) and smooth-bootstrap goodness-of-fit diagnostics with re-fitting (bottom).

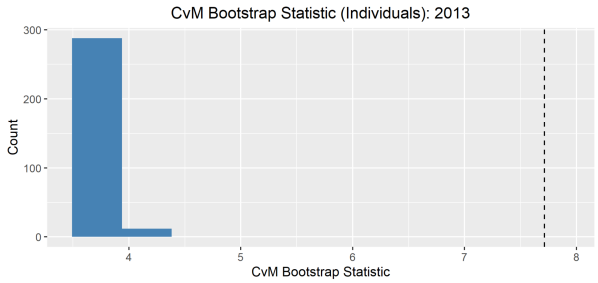
(a) Histogram and KDE overlay (individuals)



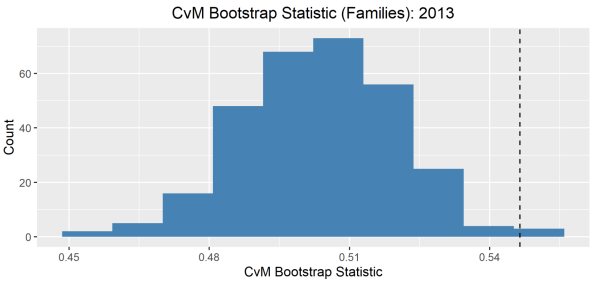
(b) Histogram and KDE overlay (families)



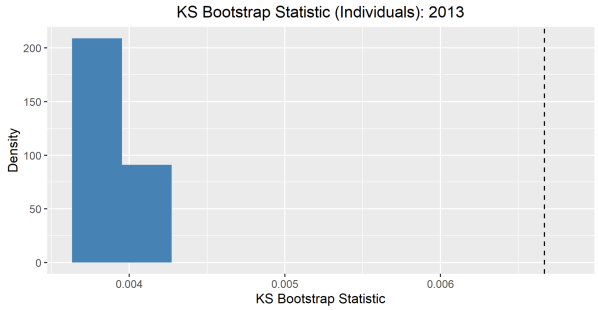
(c) CvM bootstrap (individuals)



(d) CvM bootstrap (families)



(e) KS bootstrap (individuals)



(f) KS bootstrap (families)

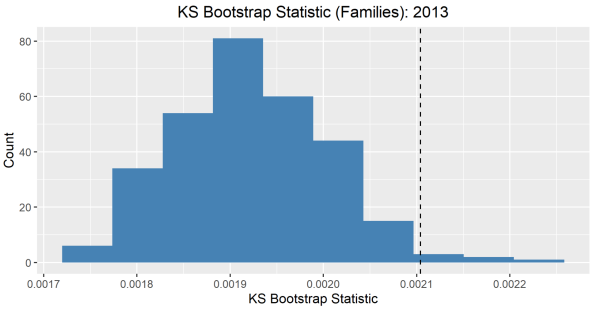
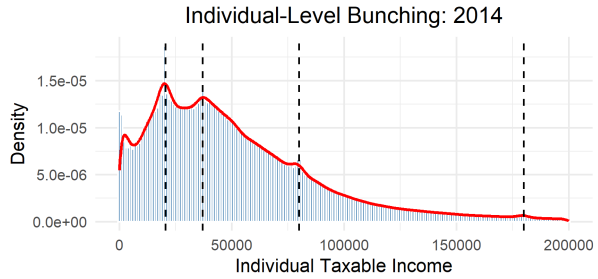
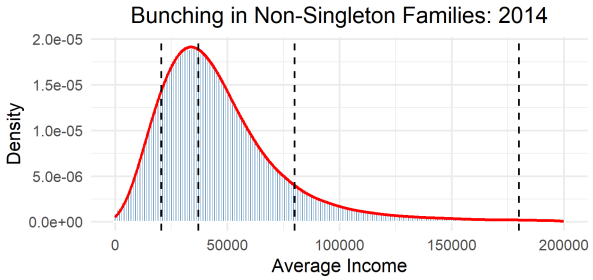


Figure 8: Year 2013: Histogram and KDE overlay with tax-threshold markers (top) and smooth-bootstrap goodness-of-fit diagnostics with re-fitting (bottom).

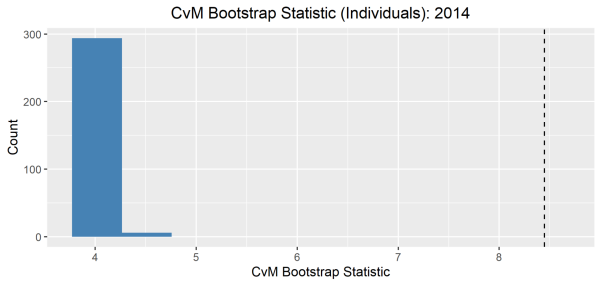
(a) Histogram and KDE overlay (individuals)



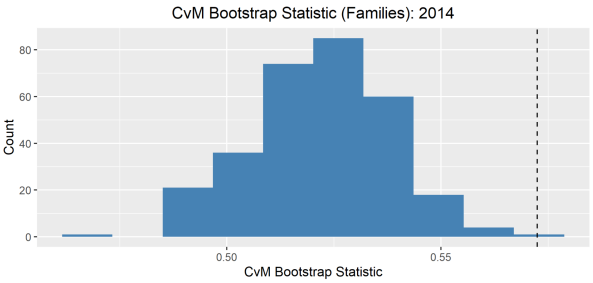
(b) Histogram and KDE overlay (families)



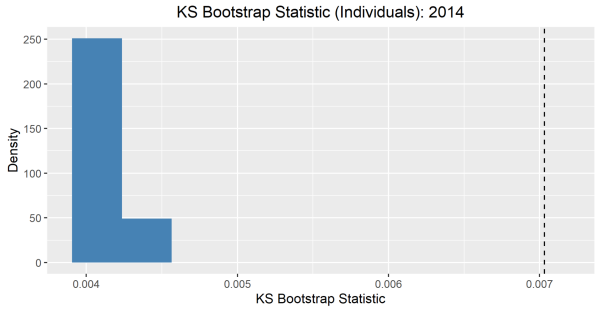
(c) CvM bootstrap (individuals)



(d) CvM bootstrap (families)



(e) KS bootstrap (individuals)



(f) KS bootstrap (families)

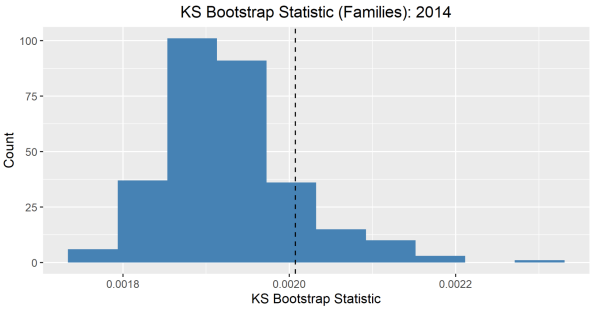
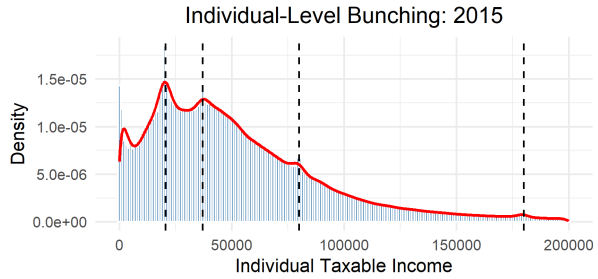
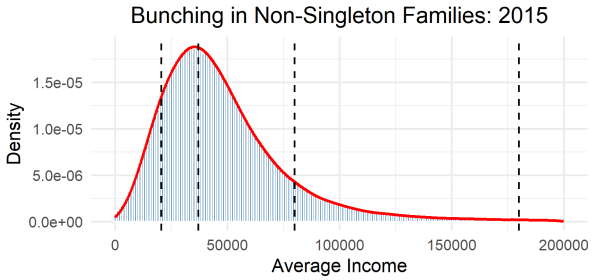


Figure 9: Year 2014: Histogram and KDE overlay with tax-threshold markers (top) and smooth-bootstrap goodness-of-fit diagnostics with re-fitting (bottom).

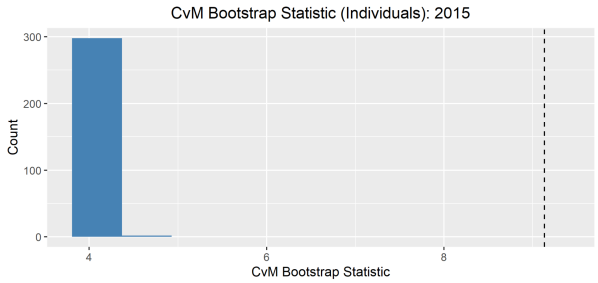
(a) Histogram and KDE overlay (individuals)



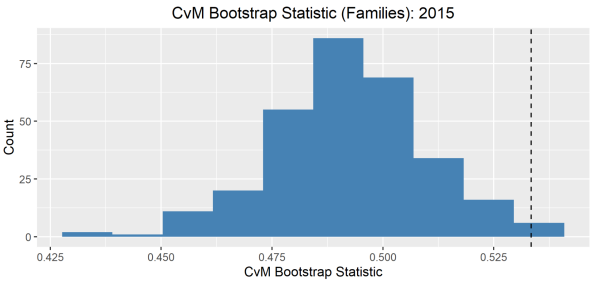
(b) Histogram and KDE overlay (families)



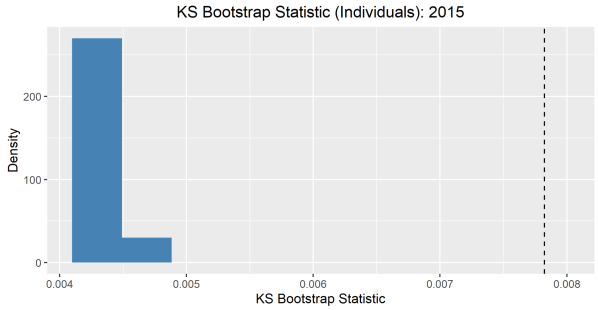
(c) CvM bootstrap (individuals)



(d) CvM bootstrap (families)



(e) KS bootstrap (individuals)



(f) KS bootstrap (families)

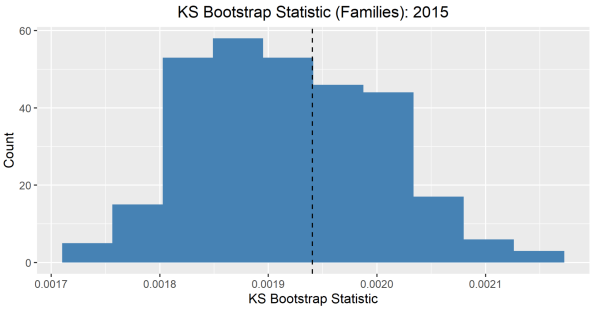
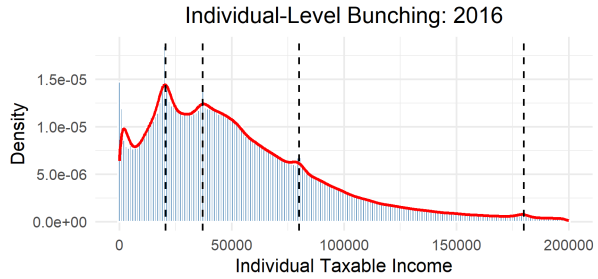
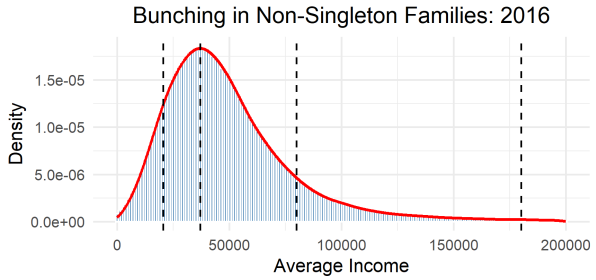


Figure 10: Year 2015: Histogram and KDE overlay with tax-threshold markers (top) and smooth-bootstrap goodness-of-fit diagnostics with re-fitting (bottom).

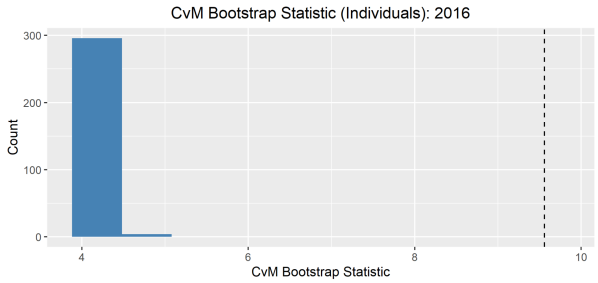
(a) Histogram and KDE overlay (individuals)



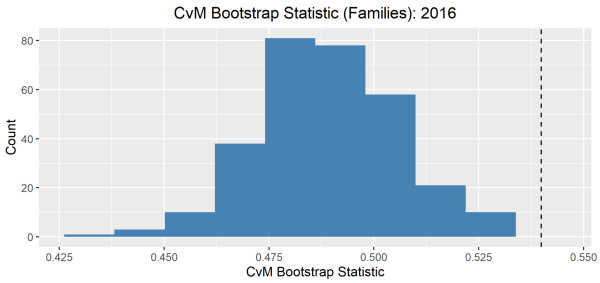
(b) Histogram and KDE overlay (families)



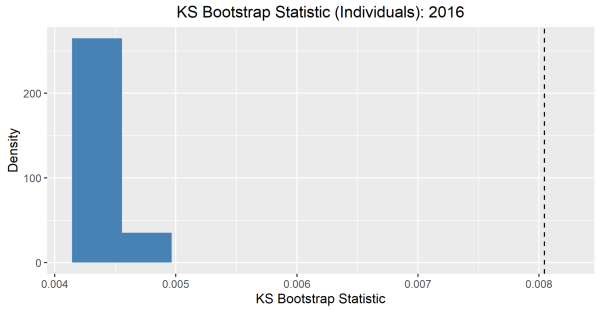
(c) CvM bootstrap (individuals)



(d) CvM bootstrap (families)



(e) KS bootstrap (individuals)



(f) KS bootstrap (families)

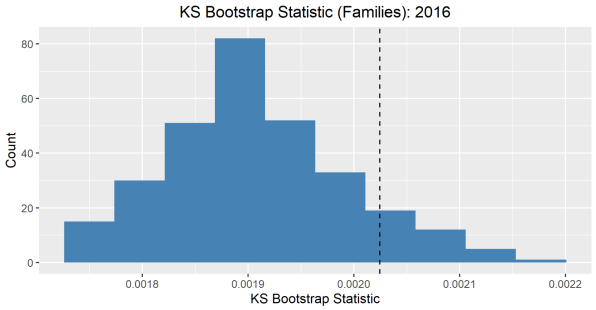
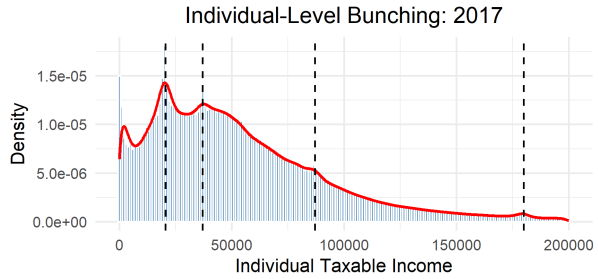
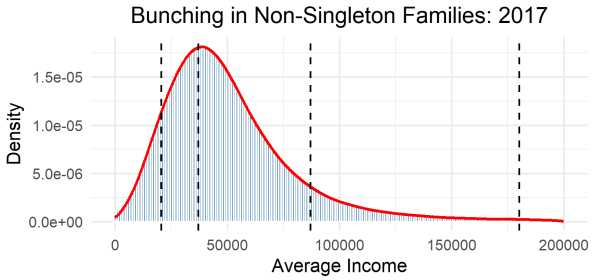


Figure 11: Year 2016: Histogram and KDE overlay with tax-threshold markers (top) and smooth-bootstrap goodness-of-fit diagnostics with re-fitting (bottom).

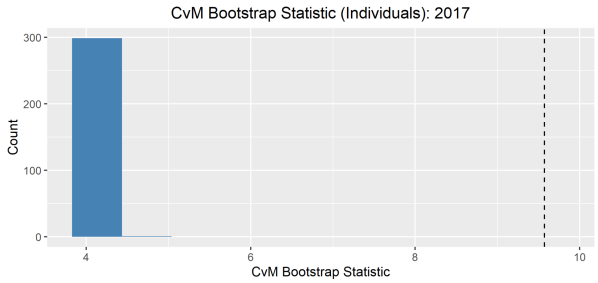
(a) Histogram and KDE overlay (individuals)



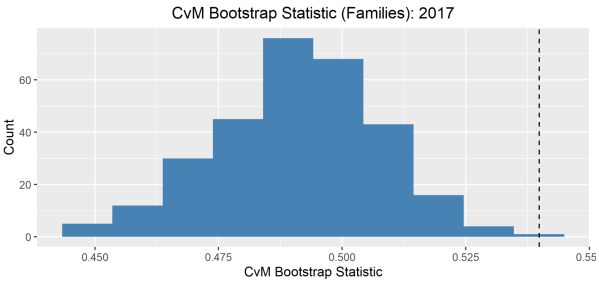
(b) Histogram and KDE overlay (families)



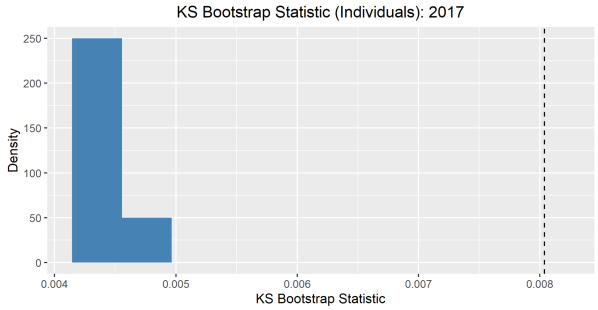
(c) CvM bootstrap (individuals)



(d) CvM bootstrap (families)



(e) KS bootstrap (individuals)



(f) KS bootstrap (families)

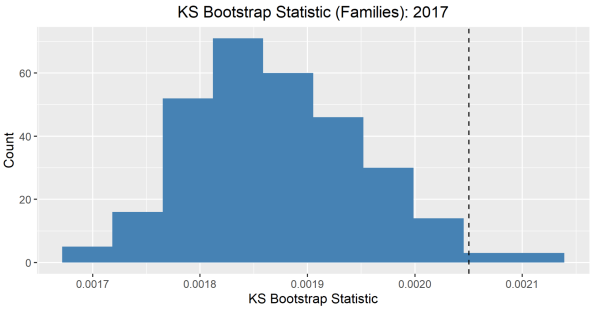
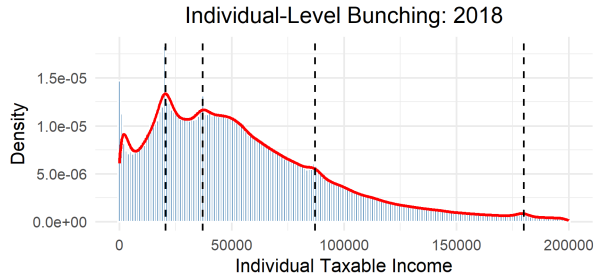
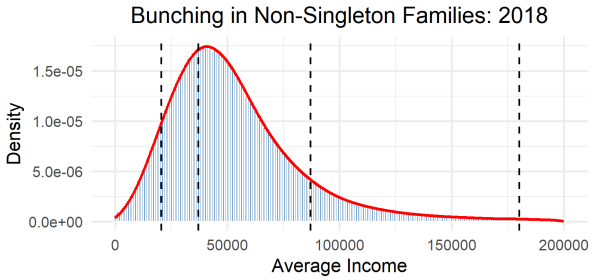


Figure 12: Year 2017: Histogram and KDE overlay with tax-threshold markers (top) and smooth-bootstrap goodness-of-fit diagnostics with re-fitting (bottom).

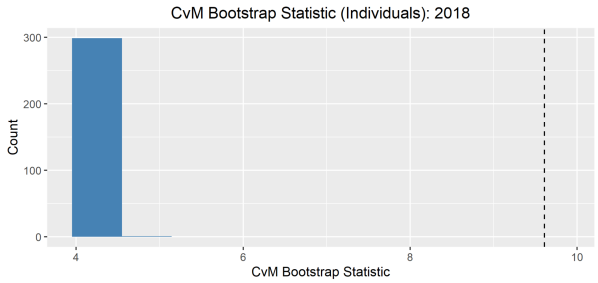
(a) Histogram and KDE overlay (individuals)



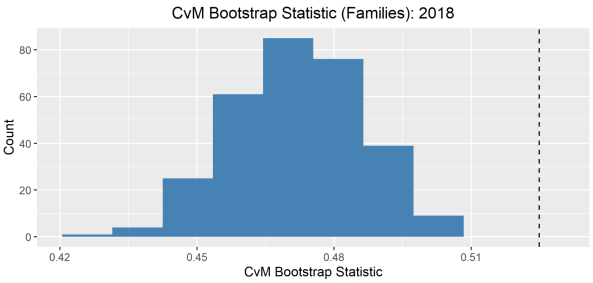
(b) Histogram and KDE overlay (families)



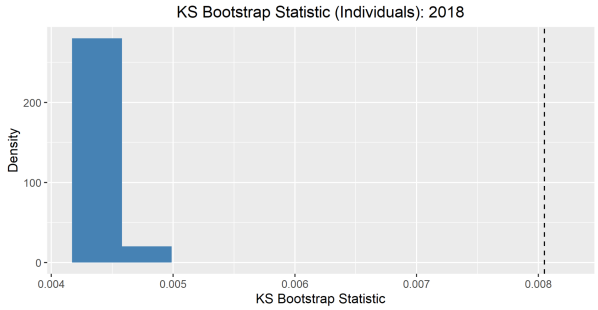
(c) CvM bootstrap (individuals)



(d) CvM bootstrap (families)



(e) KS bootstrap (individuals)



(f) KS bootstrap (families)

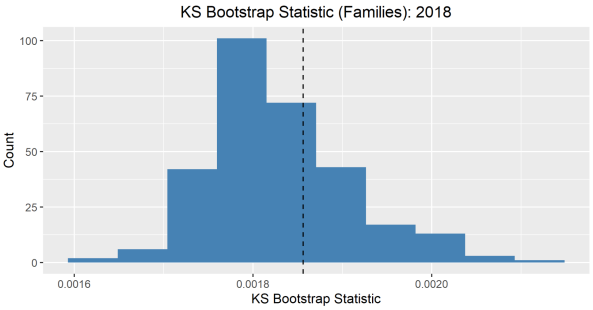
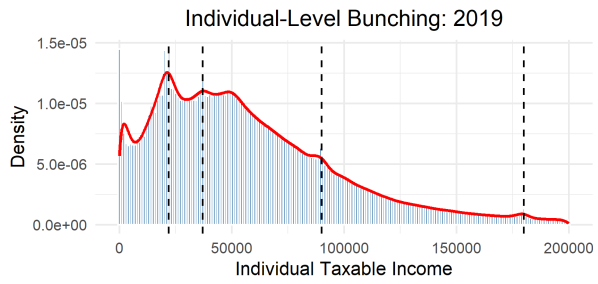
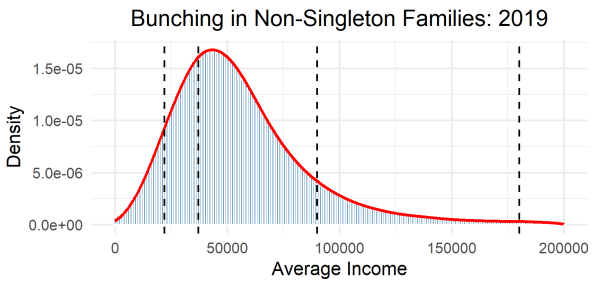


Figure 13: Year 2018: Histogram and KDE overlay with tax-threshold markers (top) and smooth-bootstrap goodness-of-fit diagnostics with re-fitting (bottom).

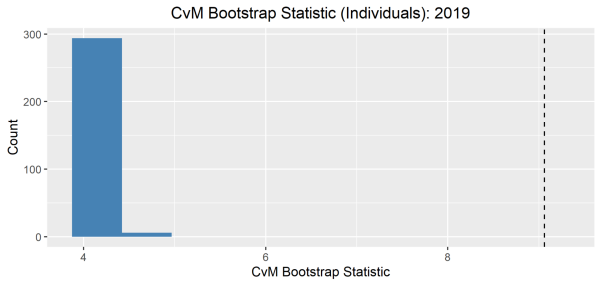
(a) Histogram and KDE overlay (individuals)



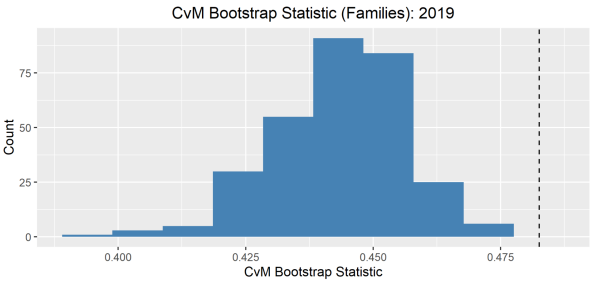
(b) Histogram and KDE overlay (families)



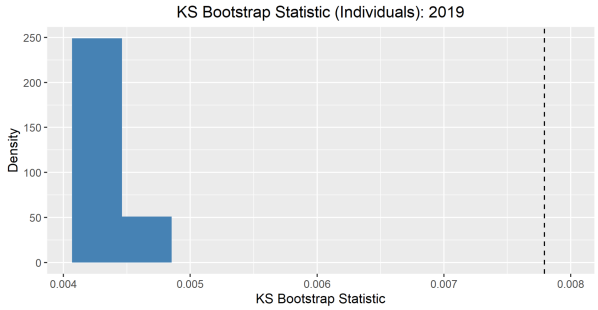
(c) CvM bootstrap (individuals)



(d) CvM bootstrap (families)



(e) KS bootstrap (individuals)



(f) KS bootstrap (families)

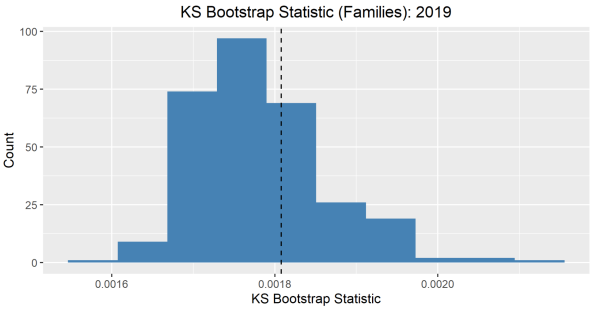
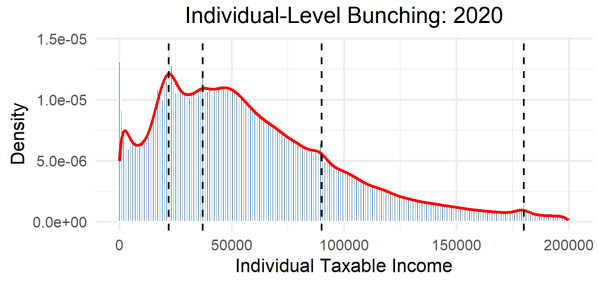
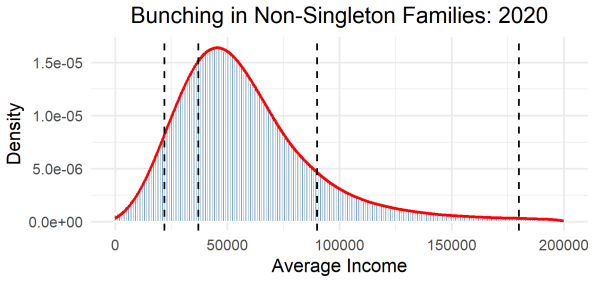


Figure 14: Year 2019: Histogram and KDE overlay with tax-threshold markers (top) and smooth-bootstrap goodness-of-fit diagnostics with re-fitting (bottom).

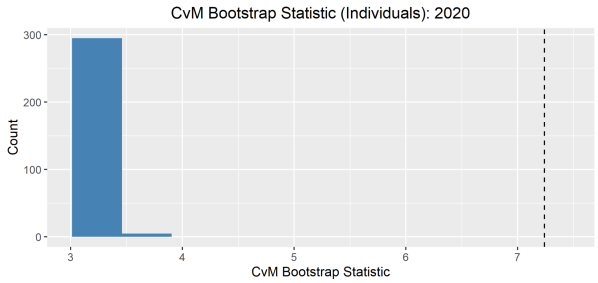
(a) Histogram and KDE overlay (individuals)



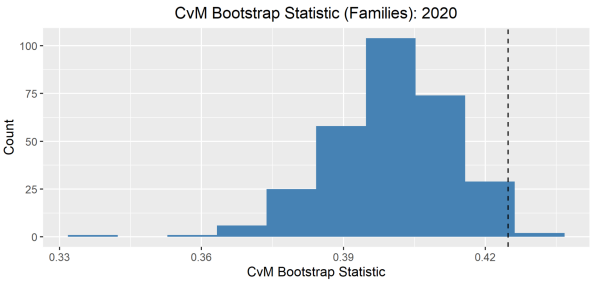
(b) Histogram and KDE overlay (families)



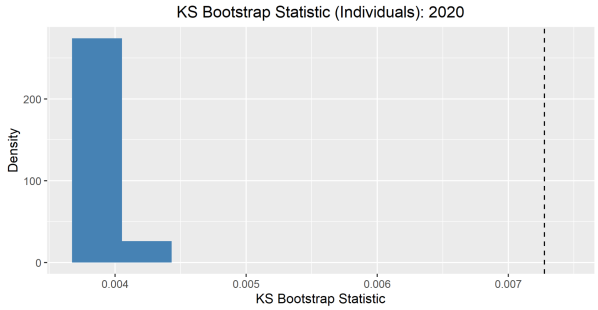
(c) CvM bootstrap (individuals)



(d) CvM bootstrap (families)



(e) KS bootstrap (individuals)



(f) KS bootstrap (families)

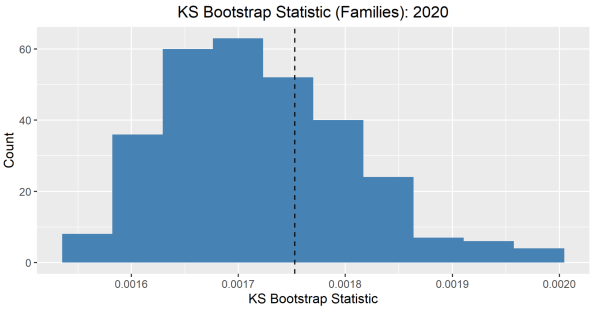
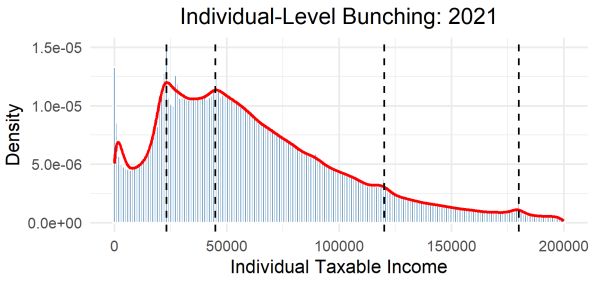
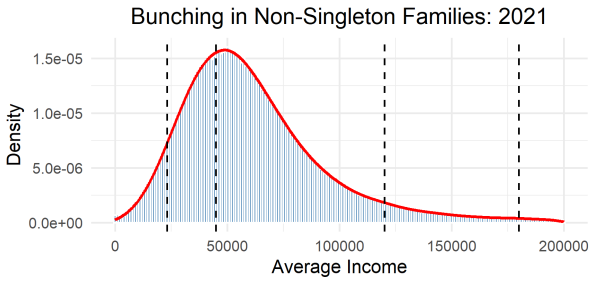


Figure 15: Year 2020: Histogram and KDE overlay with tax-threshold markers (top) and smooth-bootstrap goodness-of-fit diagnostics with re-fitting (bottom).

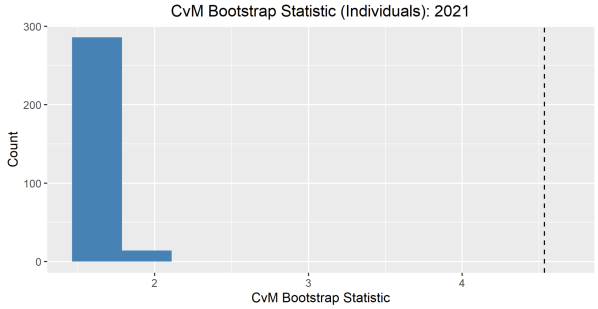
(a) Histogram and KDE overlay (individuals)



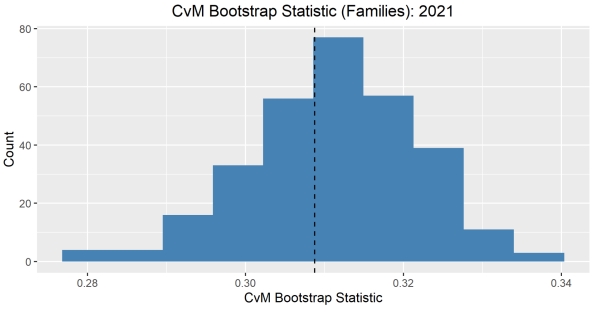
(b) Histogram and KDE overlay (families)



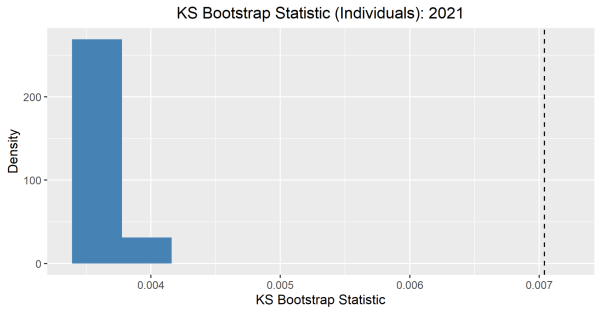
(c) CvM bootstrap (individuals)



(d) CvM bootstrap (families)



(e) KS bootstrap (individuals)



(f) KS bootstrap (families)

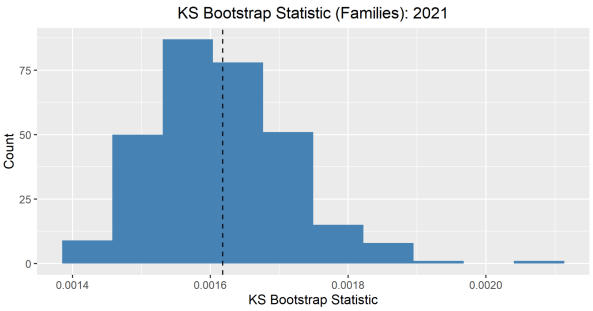
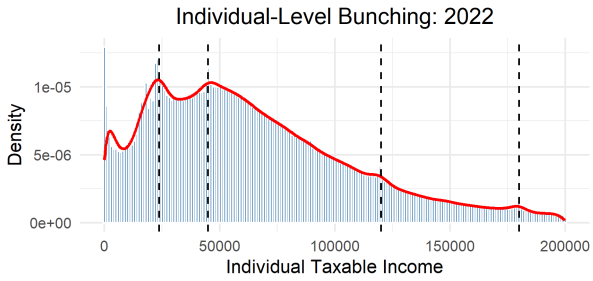
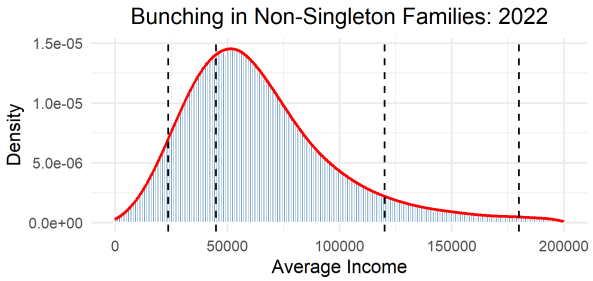


Figure 16: Year 2021: Histogram and KDE overlay with tax-threshold markers (top) and smooth-bootstrap goodness-of-fit diagnostics with re-fitting (bottom).

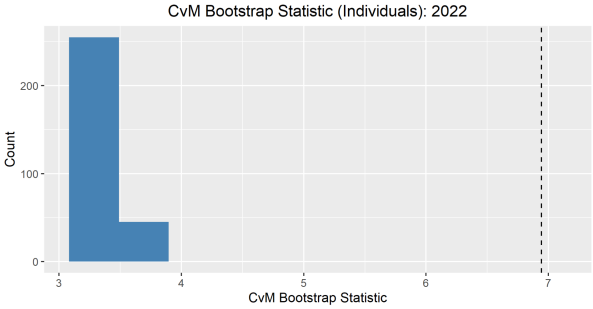
(a) Histogram and KDE overlay (individuals)



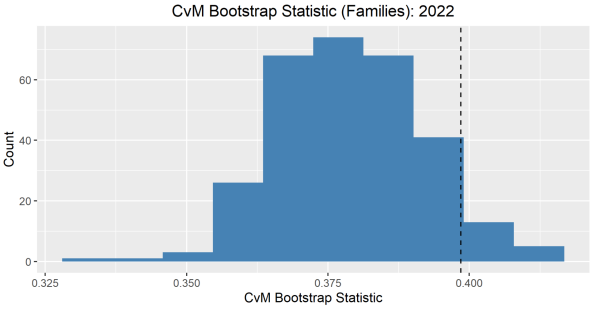
(b) Histogram and KDE overlay (families)



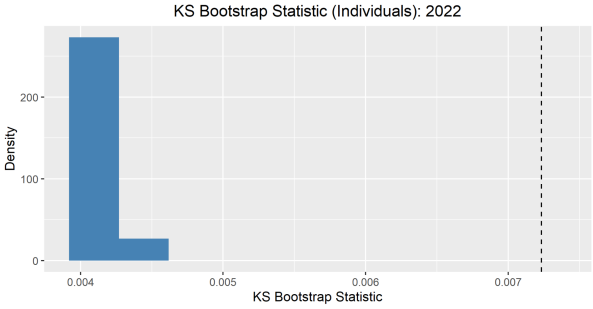
(c) CvM bootstrap (individuals)



(d) CvM bootstrap (families)



(e) KS bootstrap (individuals)



(f) KS bootstrap (families)

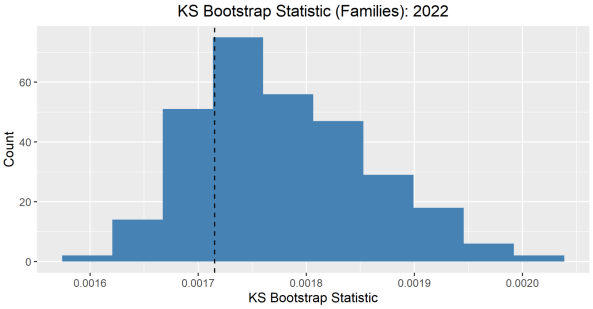


Figure 17: Year 2022: Histogram and KDE overlay with tax-threshold markers (top) and smooth-bootstrap goodness-of-fit diagnostics with re-fitting (bottom).