Oversampling and questionnaire effects on precision and coverage of wealth: the new features of the French Wealth Survey wave 2010

Abstract

The French Wealth Survey (FWS), now part of the Eurosystem Household Finances and Consumption Survey (HFCS), has been improved wave after wave since 1986. The last wave, conducted in 2010, was particularly innovative in terms of sampling design and interview process. The new procedures were designed to enhance the measure of high concentration, which is a key issue in the analysis of the distribution of wealth among households. In this article we shortly describe the various innovations for the 2010 wave of the FWS, in most cases based on the U.S. Survey of Consumer Finances experience. We then take advantage of the existing data to assess the effect of these innovations on the measure of wealth.

1 The new features of the wave 2010

The French Wealth Survey (FWS), also called Enquête Patrimoine, is a household survey that aims at measuring wealth and indebtedness of the French households. It exists since 1986 and, until now, was conducted every six years. Since 2010, this survey is included in the Eurosystem Household Finances and Consumption Survey (HFCS) that is conducted in all euro area countries. At the same time the survey benefited from various and important innovations that were supposed to increase the quality of the collected data. Assessing and sharing the results of these innovations may have implications on the methodological choices that will be made for the next waves of the survey.

1.1 Enlarging the field of the survey

Until 2010, the field of the survey excluded overseas territories (French Guyana, Reunion, French West Indies). In 2010, these territories were included in the survey, so that estimations about wealth could combine the entire territory of France.

Furthermore, not only the geographical field was extended but also the conceptual one. Indeed, in 2010 additional asset categories, for which information had not been collected previously, were also taken into account for the computation of total household wealth. These additional asset classes encompass valuables (jewelry, durables,...). This change has very strong effects on the results, since valuables and durables stand for half the assets possessed by the less wealthy, as we will see in section 2.

1.2 Oversampling of the wealthy

As already emphasised in the literature dedicated to wealth distribution among households, assets are highly concentrated in the top of the distribution of the households. For instance,

the Household Finance and Consumption Survey (HFCS), conducted for the first time in the euro area in 2010, showed that two-thirds of total household net wealth was possessed by the 20% richest European households (HFCN (2013b)). In France, the results of FWS, the French part of the HFCS, led to the same conclusion: the 10% richest households actually held about half of the net wealth, and the top percentile possessed 19% of it. As a consequence, the most heterogeneous part of the population regarding wealth is located in the top of the distribution. This can easily be illustrated by the coefficient of variation computed on every 5-percentiles of wealth (see figure 1). Following a criterion such as the Neyman one[1], we can then conclude that an efficient allocation in terms of precision will lead to an oversampling of the wealthiest households. Furthermore, certain types of assets (e.g. business assets) are possessed only by a small fraction of the population generally located in the top of the distribution.

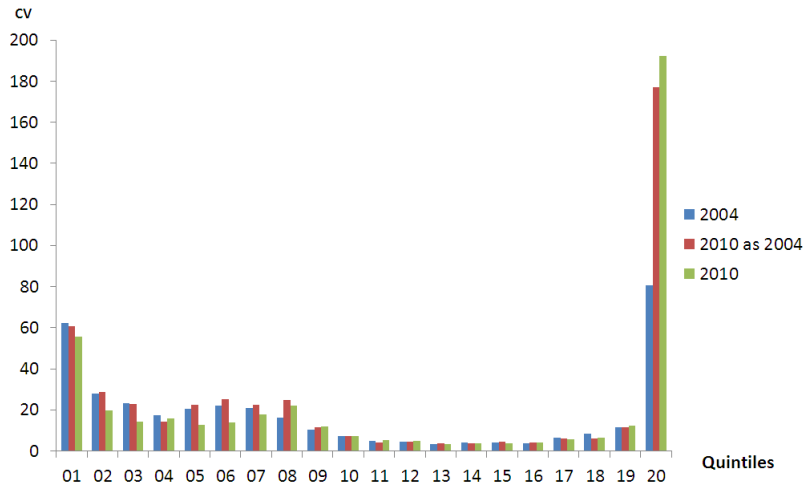

Figure 1: Coefficient of variation for FWS waves 2010 and 2004

Note: The "2010 as 2004" CV designates the coefficient of variation computed on the 2010 data, with the methodology used in 2004. For more information, see the section 1.4

The oversampling has been successfully implemented in surveys such as the U.S. Survey of Consumer Finances. For example, Kennickell (2008) describes very precisely the methodology used to oversample the richest part of the population. Indeed he explains how important the wealthiest 1 percent of households is for the estimation of total household wealth. Following the U.S. example, the Eurosystem Household Finances and Consumption Network (HFCN), in charge of the HFCS, decided to adopt the same objective. This was done with various methodologies. For an overview of these different methodologies, see HFCN (2013a).

We focus more precisely on the French case, which has adopted the oversampling strategy in 2010. Even if sampling design for the four previous waves was taking into account wealth information, it rather focuses on categories in the population that were more likely to detain specific assets, in particular self-employed and elderly persons and did not implement such an ambitious oversampling strategy. In 2010, the sampling base was built thanks to administrative data, in particular housing tax data combined with income tax data. Thus the register used for 2010 contained a wide range of information about assets and income. It was then possible to focus on the households that were more likely to belong at the top of the distribution. The

[1] see Cochran (1977)

oversampling part combined 3,000 households, while 17,000 households were selected for the "standard" part of the population. To compare the oversampling between 2004, year for the previous wave of the FWS, and 2010, we then use the effective oversampling rate, that can be defined as follows, focusing on the last decile:

$$r_\alpha = \frac{S(100-\alpha) - \alpha/100}{\alpha}$$

where $S(100-\alpha)$ denotes the share of the households belonging to the highest α-percentile in the sample. Thus the oversampling in 2004 was quite limited: only 13.7% of the sample belong to the last decile of wealth, which brings the effective oversampling ratio to 37. In 2010, the effective oversampling ratio was 129 (see HFCN (2013a)), which concurs to validate the oversampling strategy chosen for this wave. Put in other words, it means that the households belonging to the 10% wealthiest represent about 22.9% of the sample, more than twice their share in the population. The oversampling strategy enables accurate analysis on the very top of the distribution: in 2004, with an effective oversampling ratio of 37, the sample contained only about one hundred households in the highest percentile, limiting *de facto* the analysis on this group. In 2010, with an effective oversampling ratio of 438 and more than 800 households in the sample, analysis combining not only the level of wealth but also portfolio choices and sociodemographic characteristics was possible.

1.3 Questionnaire strategy

Since an important part of the sample was likely to possess high amounts of assets, the questionnaire was modified to take into account this possibility. Usually, asking questions about assets can be demanding, and respondents feel difficulties or even are reluctant to give precise amounts for the assets they declare to the interviewer. This issue has been documented in the literature; for the French Wealth Survey, see for example Arrondel, Guillaumat-Tailliet, and Verger (1996). The main reasons for underreporting or not reporting at all can be very different from one household to the other. Risk-averse households would declare the value of their financial assets with extreme caution, contributing to the mismatch between global financial amounts measured in the survey and financial accounts. Other households would consider wealth as part of their private life and thus would refuse to provide precise amounts. With no doubt, social control is at work here, since the main residence, whose value can be assessed by the interviewer, is the least underreported asset.

Given the difficulties to collect precise amounts, a classical strategy is to ask households questions about amounts with brackets. Since they can provide fuzzy information, households are less reluctant to answer. On the one hand, they can give an answer that take into account all the uncertainty related to the market value of their assets. On the other hand, the less precise amounts they can give help in the process of anonymisation of the data and lower the exposure risk, with definitively less damages for the measure than no response at all. Furthermore, the respondent uses generally cards to answer questions about amounts: from this point of view, using cards and designating the bracket to the interviewer with its letter rather than with its lower and upper bounds contributes to reassure the respondent.

Once accepted the principle of answers in brackets, the following issue is the thickness of these brackets. Too many brackets would complicate the answer of the household; too few would make it more inaccurate. According to the different types of assets, the number of brackets does not exceed 12, with a top bracket that includes all amounts above EUR 450,000. If such an amount can appear to be quite high for a given asset, it becomes insufficient when focusing on the top of distribution. Consistently with the oversampling strategy, it became obvious that an

Response	Proportion of the declared financial assets	Share in the global value of financial assets
Amount declared as a number	34.0%	43.5%
Amount declared as a bracket, value over EUR 450,000	0.4%	8.7%
DK-DA	*3.4%*	*2.2%*
EUR 450,000 up to 1 million	*55.6%*	*26.0%*
EUR 1 up to 2 millions	*25.4%*	*24.4%*
EUR 2 up to 5 millions	*10.8%*	*18.5%*
EUR 5 up to 10 millions	*2.2%*	*14.2%*
EUR 10+ millions	*2.6%*	*14.7%*

Table 1: Response pattern for the amounts of financial assets in FWS 2010

increase of the number of proposed brackets was necessary. To do so, a two-step strategy was followed: respondents were offered the same set of brackets than in 2004, but those choosing the top bracket were given a new set of brackets, with amounts from EUR 450,000 up to EUR 10,000,000 and more. As shown in table 1, this contributes to improve substantially the answer of the households belonging to the top of the distribution. Indeed 94% of their assets were provided with an upper bound that would have remained unknown with the 2004 questionnaire. Enhancing the accuracy of the estimated amounts for the top of the distribution is essential, since this 0.4% of assets over EUR 450,000 stands for 8.7% of the total value of financial assets estimated with the survey. As a consequence, the collected amounts should be more accurate.

1.4 Imputation strategy

One key issue in surveys aiming at measuring wealth is to estimate a global amount of assets. Indeed, since most of the households provide only amounts in brackets, it is difficult to sum all assets to obtain a global amount. One solution could be summing both lower and upper bounds of assets to obtain a minimum and a maximum, but such a computation would lead to a too inaccurate measure of the household's wealth. The main solution usually used for this kind of survey is to impute an amount unsually constrained to remain between the lower and the upper bounds, which then allows easily to compute global amounts. To do so, the general idea is to estimate an econometric model on the data, take the predicted value and add a residual that makes the imputed value fit with the answer of the household. This "simulated residual" method has been for example described by Gourieroux, Montfort, Renault, and Trognon (1987). In our case, given an asset i, we need to estimate a censored model, taking into account that the real amount y, here unknown, is lower than M and greater than m. Thus we make the assumption that y_i depends on a set of covariates, X_i. We want to estimate the following model:

$$log(y_i) = \beta X_i + u_i$$

Making the assumption that the error term, u_i, follows a normal law, we then can write the likelihood of the model:

$$L_i = \Phi(\frac{M - \beta X_i}{\sigma}) - \Phi(\frac{m - \beta X_i}{\sigma})$$

Once $\hat{\beta}$ is estimated, it is possible to compute a value $\hat{y_i}$ for y_i, adding to the predicted part βX_i a residual u_i that respects a set of constraints, in particular the fact that $m \leq \beta X_i + u_i \leq M$.

It is possible to generate u_i thanks to an acceptation-rejection algorithm or with a Gibbs sampling algorithm. However, the quality of the model is here essential: if the βX_i are far from the declared brackets, this would impose a potentially high number of iterations before finding the proper u_i. Furthermore, the estimated model gives a valuable answer for the imputation process to the following question: is the amount in average at the top, in the middle or at the bottom of the range?

This is the reason why concerns should focus on the covariates X_i. The quality of the predictors is a key point here, which has been taking into account for the 2010 wave of the FWS. A wide range of covariates from the register data have been added in the models. Moreover, information about households' wealth have been added as constraints for the generation step of the residuals. This has of course an effect on the estimations obtained thanks to the data.

2 Assessing the effect of innovations on the survey

2.1 Enlargement of the field

The inclusion of the overseas territories has, as expected, an effect on the estimation of wealth. However, the population living in overseas territories stands for less than 3% of the total population living in France. Their inclusion in the sample was not expected to increase significantly the global amount of wealth. Indeed, the assets possessed by overseas population represent about 1.6% of the total of assets in France.

On the other hand, the incorporation of assets such as valuables and durables to the survey has led to a relative increase of the wealth for the bottom of the distribution (figure 2). Indeed, these assets stand for a meaningful fraction of the wealth of the least wealthy households: for example, half of the assets owned by the 10% least wealthy households belong to this category (see Lamarche and Salembier (2012)).

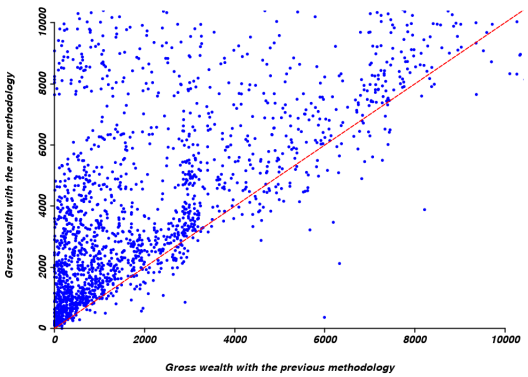

Figure 2: Effect on the gross wealth for the household below EUR 10,000 with the previous methodology

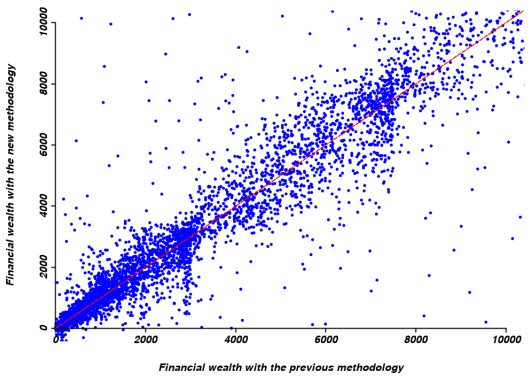

Figure 3: Effect on the financial wealth for the household below EUR 10,000 with the previous methodology

2.2 Effect of information collection

To assess the effect of the new methodology on the wealth amount estimate, we need a counterfactual that reproduces the previous methodology. What we call here as "new methodology"

combines both the new set of brackets and the improved models for imputation. A new dataset is computed by suppressing the extra information specific to 2010: the brackets are re-computed as if the households were given the those used in 2004 and the models for imputation exclude variables that were not available in 2004. This demanding procedure implies to reinitiate all the process of statistical production but is necessary for a survey repeated across time as the FWS: indeed the survey exists since 1986 and it is essential to keep producing not only cross-sectional analysis but also comparisons over time of the repartition of wealth across households. Thus all the figures describing the evolution of wealth between 2004 or 1998 and 2010 have been computed thanks to this dataset providing comparable estimates.

To assess the effect of the new methodology on the estimations, we compare results between the wealth distribution obtained with this methodology and the one that would have been computed with the previous methodology. It enables to identify which part of the distribution is mainly concerned by the innovations and identify more accurately the effect of the new methodological choices.

Estimators	Gross wealth	Financial wealth	Real estate assets
Mean	10%	10%	3%
D1	215%	-4%	-
D2	137%	-2%	-
D3	36%	2%	-
D4	6%	0%	-
D5	6%	1%	1%
D6	8%	3%	2%
D7	7%	2%	3%
D8	7%	4%	3%
D9	7%	5%	2%
P99	13%	11%	3%

Figure 4: Gap between previous and new methodologies

As shown in table 4, the new methodology has had an effect on the overall distribution, with slight differences regarding the position in the distribution. Indeed, the first decile of gross wealth is multiplied by more than 3 with the new methodology. However, as also illustrated by the figure 2, the extension of the conceptual field, including now also valuables and durables, explains mainly this impressive increase for the bottom of the distribution. If we now focus on financial assets, for which there is no modification in the definition, but only changes in the data collection, the effect on the bottom of the distribution is far less outstanding (see figure 3). On the contrary, the two first deciles are decreasing with the new methodology, whereas the others are increasing. The strongest effect concerns the top of the distribution, without any surprise: the last percentile of financial wealth increases by 11% thanks to the new methodology, driving the increase of the last percentile for gross wealth. Thus the introduction of a new set of brackets and the use of register variables in the imputation models has a significant effect on the top of distribution, consistently with the oversampling strategy, while low or no effect on the rest of the distribution. Furthermore, the effect is quite negligible when focusing on real estate assets: here the brackets have not been modified since households set themselves a minimal and a maximal amount for the value of their real estate assets. The only modification concerns the models that include new variables and new constraints for a part of the households for whom the amount of real estate assets is available in the registers.

2.3 Effect of the oversampling

In the Neyman approach, oversampling aims at enhancing the accuracy of the survey by over-representating (beyond proportionality) in the sample the most heterogeneous groups of the population. Such a strategy should have then an effect on the precision of the estimators computed on the data of the survey. A first assessment consists of the comparison between the precision for given estimators obtained in 2004 and those obtained in 2010. The estimators that we retained are the most commonly used ones, namely the gross wealth mean and deciles. However, such a comparison could be misleading since not only the estimators are more accurately measured but also their precision. Indeed, the precision of an estimator depends generally on the term S^2, that can be assimilated to the variance of the variable. If this variance is more accurately measured, the estimator can paradoxically be considered as less accurate. This phenomenon can explain partly the results shown in figure 1, for which we observe a strong increase in the coefficient of variation in the top of the distribution. This increase could either be explained by an increase if the heterogeneity of this part of the population or by the improvement of the measure.

To tackle this kind of issue, one would want to build a counterfactual to assess the gain in terms of precision due to oversampling. Such a counterfactual can be approximated thanks to the design effects that are already themselves a kind of counterfactual. Indeed the idea behind the design effect is the comparison between the precision obtained for the sampling design and the precision that would have been obtained with a Simple Random Sample (SRS), as shown in the formula of the design effect:

$$DEFF = \frac{\hat{V}}{\hat{V}_{SRS}}$$

Thus a design effect below 1 will show the efficiency of the sampling design, while a design effect over 1 would lead to the conclusion that, in terms of accuracy, a simpler design sampling would have been preferable. Results shown in table 2 prove that the oversampling strategy is efficient, as expected, on the top of the distribution: the design effect associated to the estimator of the last percentile of gross wealth, for example, is below 1. However, the rest of the distribution does not benefit from this strategy, since the design effects for the other quantiles are more than 1. Median of gross wealth has a design effect of 1.41, meaning that the uncertainty associated to the estimation of this indicator would have been reduced by 30% if using a SRS. This indicator is although estimated accurately: the coefficient of variation of its estimator is 2%, which represents an error of EUR 3,000. The oversampling strategy is the result of a trade-off between the precision of various estimators. Indeed it allows enhancing the estimation of the top of the distribution and the mean, that is highly driven by the top, while fading the rest of the distribution.

As the survey relies on a complex sampling design, the computation for the precision of the estimators is quite demanding. The results presented here are given by an analytical procedure, based on a recursive method (see Durbin (1953), Raj (1966), Rao and Lanke (1984)) and combined with a linearization method for non-linear estimators (such as quantiles, for more information see Deville (1999)). Furthermore, replicate weights have been computed in the FWS as core variables for the HFCS. Thus it is possible for any user to re-compute this estimators thanks to the provided data.

Finally, it would be informative to assess the gain in precision obtained thanks to oversampling. Indeed, the design effects do not provide complete information because they also take into account the stratification built for the sampling design and the different degrees of sampling. One way to approximate the needed counterfactual is to compute design effects not only for the whole sample but also for each stratum. Then it is possible to recombine the stratum allocation without any oversampling strategy and, taking into account the response rate for each stratum,

Estimators	Gross wealth	Net wealth	Indebtedness	Real estate assets	Financial wealth
Mean	0.65	0.69	0.74	0.74	0.31
P99	0.72	0.73	1.10	0.83	0.75
P95	1.23	1.21	1.38	1.56	1.28
P90	1.32	1.36	1.49	1.70	1.47
Q3	1.50	1.41	1.43	1.72	1.56
Median	1.41	1.35	1.47	1.51	1.58
Q1	1.68	1.72	1.47	1.37	1.79
P10	2.00	1.91	1.47	1.37	1.88
P5	1.86	1.97	1.47	1.37	1.88

Table 2: Design effects for the most commonly used indicators in the FWS 2010

compute a variance as it would have been expected with the stratification but without oversampling. Thus, taking advantage of the additive property of the variance for a stratified sample, we can use the following formula:

$$V(Y) = \sum_{h=1}^{H} (\frac{N_h}{N})^2 (1 - \frac{n_h}{N_h}) \frac{S^2_{Y,h}}{n_h} DEFF_h$$

where N, N_h, n_h and $DEFF_h$ respectively denote the size of the population, the size of the population in stratum h, the size of the sample in stratum h and the design effect estimated in the stratum h. Thanks to this formula, it is possible to assess the precision of a given estimator for different allocations of the sample between strata.

Then the initial sample would have included only 400 households belonging to the top of the distribution, against 3,000 with the oversampling strategy. Thanks to the described back-of-the-envelope calculation, we find that the variance associated to the mean gross wealth would have increased by 48% without oversampling. This implies an increase of the uncertainty by EUR 1,200. It also gives a tool to arbitrate between increasing costs associated to oversampling and gain in precision for given estimators. Indeed, the response rates are lower in the oversampled part of the population and the duration of the questionnaire increases, due to the fact that these households have more assets to describe. In that sense, information is more costly to collect among these households.

3 What about underreporting?

As previously emphasised, the survey suffers since the beginning from an underreporting phenomenon that turns out to be common to the entire euro area. Indeed, when comparing global amounts measured in the national accounts and those given by the survey, we find that mean amounts in the survey are far below those expected. This issue has already been discussed, for the French case, by Arrondel, Guillaumat-Tailliet, and Verger (1996). In 2004, the coverage rate for gross wealth was 58%. The innovations introduced in 2010 aim in particular at improving this coverage rate, which would be a positive indicator for the quality of the survey. To this point of view, the various innovations have contributed to the enhancement of the quality of the survey. Indeed the coverage rate reaches 69% in 2010. Furthermore, this kind of comparison is not completely relevant, because of wide differences in field and concepts between national accounts and the survey. An exercise of so-called "calibration" on the national accounts has been run and

compared with the previous one made on the 2004 wave (see Cordier and Girardot (2007) for more information about the method and the results in 2004), which shows interesting decreases for the calibration coefficients. In particular, financial assets, for which the lowest coverage rates are usually observed, are better recovered. For example, life insurance in the 2004 wave stands for only 28% of the national accounts. This figure in 2010 was higher and reaches 42%.

4 Conclusion

Summing up these results we conclude two things:

1. First, the innovations that have been introduced are efficient, since they have contributed to an increase of the assets covered by the survey. In particular, field extensions have an important effect. For example, taking into account valuables and durables has enabled to increase the mean gross wealth by 4.7%, the inclusion of overseas territories by 1.6%. The rest of the increase is due to oversampling and improvement of both questionnaire and imputation models, without any possibility to disentangle the different effects.

2. Second, after taking into account as far as possible for differences of field and concepts between the survey and the national accounts (for a complete description see for example HFCN (2013a)), all these innovations seem not to be sufficient to solve the underreporting phenomenon. Other strategies could be investigated, such as the more extensive use of registers to lighten the load of interrogation for the households. Indeed, participation and amounts could be retrieved thanks to registers (if, of course, available) but this kind of solution needs to be tested before all to assess its impact on the wealth estimate. This would among other things permit to shorten substantially the questionnaire and focus it more deeply on qualitative questions, such as anticipations that are key variables in the analysis of wealth accumulation.

References

Arrondel, L., F. Guillaumat-Tailliet, and D. Verger (1996): "Montants du patrimoine et des actifs: qualité et représentativité des déclarations des ménages," *Economie et statistiques*, 296(1), 145–164.

Cochran, W. G. (1977): "Sampling techniques," *New York: John Wiley and Sons*.

Cordier, M., and P. Girardot (2007): "Comparaison et recalage des montants de l'enquête Patrimoine sur la comptabilité nationale," *Document de Travail INSEE*.

Deville, J. C. (1999): "Variance estimation for complex statistics and estimators: Linearization and residual techniques," *Survey methodology*, 25(2), 193–204.

Durbin, J. (1953): "Some results in sampling theory when the units are selected with unequal probabilities," *Journal of the Royal Statistical Society: Series B (Methodological)*, 15(2), 262–269.

Gourieroux, C., A. Montfort, E. Renault, and A. Trognon (1987): "Simulated residuals," *Journal of Econometrics*, 34(1), 201–252.

HFCN (2013a): "HFCS Methodological Report for the First wave," .

——— (2013b): "HFCS Report on the Results from the First wave," .

Kennickell, A. B. (2008): "The Role of Over-sampling in the Survey of Consumer Finances," *Irving Fisher Committee Bulletin*, 28, 403–408.

Lamarche, P., and L. Salembier (2012): "Les déterminants du patrimoine : facteurs personnels et conjoncturels," *Insee Références Les revenus et le patrimoine des ménages*.

Raj, D. (1966): "Some remarks on a simple procedure of sampling without replacement," *Journal of the American Statistical Association*, 61(314), 391–396.

Rao, J., and J. Lanke (1984): "Simplified unbiased variance estimation for multistage designs," *Biometrika*, 71(2), 387–395.