

The winning combo

Introducing a passive digital measurement in an NRS

Nicolas COUR – Managing Director, ACPM

Gilbert SAINT JOANIS – Chief Research Officer, ACPM

INTRODUCTION

There is no longer any doubt about the importance of digital reading to press brands. The constant growth in the possession of devices, mobile devices in particular, and in the usage of these devices has resulted in a situation where the digital readership is dominant over the print readership and, within the digital, where reading is pre-eminently done via smartphone rather than via a fixed internet connection.

As such, it is essential to look at future audience measurement not as the juxtaposition of figures for separate media platforms but as an integrated tool capable of evaluating the performance of press brands as a whole.

Having made this observation, implementation is not as simple as one might hope. The appropriate methodology must first be identified and then put into practice. Above all, this must be done without distorting the value of the measurement.

Based on the principle that, in the short term, passive measurement of press readership is difficult to achieve, as we discussed in our PDRF 2015 paper, we needed to implement the measurement of digital readership in those same individuals who initially responded to the press questionnaire. We have now done this. We drew on the experience gained from our audience survey on Executives and members of high-income households. This survey was particularly useful to this trial as we do not have, for the French market, a digital audience survey for targets as detailed as those addressed in the Premium survey.

In this paper, therefore, we will be presenting the measurement methods we have put in place. We will then explore the key technical points for success and the main methodological lessons learned. Lastly, we will conduct a comparative analysis of the digital audience results produced using a declarative measurement system and those produced using a passive system, both performed on a single sample.

I- PASSIVE MEASUREMENT TECHNIQUE FOR DIGITAL AUDIENCES

Digital audiences, particularly those of press brands, can be formed of two complementary access types: websites and applications (apps). In order to conduct a complete measurement, therefore, we needed to implement two measurement processes tailored to each type of contact.

a. Website audience measurement

The audience measurement system chosen for websites is a Site Centric measurement. The installation of a meter, in addition to its development cost, would have generated much greater resistance due to its more invasive nature. Website audiences are therefore measured using a tag/cookie system. However, unlike normal Site Centric measurement, the cookie used in this survey carries a unique identification number allowing us to track the individual internet usage of each of the recruited panellists.

As such, at each visit to any of the websites participating in the survey, the cookie number is identified and added to the audience database. This technique combines the flexibility of Site Centric measurement with the capacity to measure at individual level.

Another advantage of this technique is that cookies can be installed on all types of devices and in any environment. Cookies are compatible with all browsers and can be just as easily used on a computer, tablet or mobile phone. Lastly, in the vast majority of cases, cookies are not blocked by the IT systems of large companies, unlike the use of a meter which, at the very least, requires that the user is an administrator of their workstation.

This final point was vital for us. The Premium survey, into which we implemented this type of measurement, is targeted at business executives in particular. In addition, it measures, among other things, the audience of professional publications and their websites. It was therefore essential to be able to collect the audience information for these panellists.

b. Application audience measurement

Measuring the audience figures generated by apps is more difficult. Due to the fact that apps run independently, we had to find a way to ensure that the measurement system stored information when the app was used offline. The choice we made was to use the advertising identifier available in both Apple and Android environments. Due to the low penetration rate of other mobile operating systems, we decided only to develop the tool for these two systems.

From a technical perspective, measurement takes place in two stages:

- Initial identification of the individual: we developed a dedicated app for the survey. We asked the panellists to install this app on their various mobile devices. When installed, this app sends the device's advertising identifier to our servers.
- Visit measuring: we then introduced an extra line of code (called an SDK), which activates at each connection to the app, into the apps of the publishers that wanted to be measured in the Premium survey. This SDK also sends our servers the identifier of the connected device. By consistently checking these against the restrictive list of panellist IDs we can track the panellists' contacts.

II- THE RESEARCH TOOL

a. The measurement tool

i. Use of a sample as a panel

Before discussing audience results, we need to make clear that our survey is conducted on a group of individuals who possess some characteristics of a sample but who also, in some ways, resemble a panel. We have nick-named this situation the use of a "pseudo panel".

We therefore have a panel with its own characteristics, including the option to measure digital reading over the short or long term for each individual on the panel. However, the fact that this panel is recruited in the context of a survey means that the panel is progressively extended during the survey period. As such, for some of the individuals, recruited in the early days of the survey, we will have several months of passive measurements. At the other end of the scale, for the last individuals recruited we may potentially have no passive measurements, as the survey comes to an end. In order to have a minimum amount of data available, we decided to continue the passive measurement for one month after the end of the recruitment period.

However, the consequence of this type of recruitment and implementation of passive measurement is that we have a "tapered" panel, with a very few individuals for whom measurements have been taken over a significant time scale and with less and less information available for the subsequent individuals recruited as the panel was built up.

Naturally, this situation requires a missing data injection phase in order to compensate for the imbalances associated with the progressive recruitment of the panel.

ii. Identification of the individuals and right to privacy

One of the key points of this recruitment phase is to make sure that there is complete concordance between the various elements collected concerning one individual. We are going to have:

- Information from the file used to contact interviewees
- Information provided over the telephone by the interviewee him/herself
- Information provided in the main questionnaire
- All digital audience data associated with this same individual

We therefore have to generate a unique identifier and make sure that it is present at every stage of the survey.

Aside from the technical challenge this poses, it also raises the question of data protection and confidentiality. We must not only inform the respondent about the various parts comprising the survey, but also collect their agreement to take part in each stage of the survey.

This agreement to participate is naturally an obstacle to optimising the volume of respondents taking part. It also restricts the time during which the information is collected. Providing a clear and honest description of the measurement tool forces us to clarify its extent and limitations to the respondents, particularly in terms of time scale.

We therefore know in advance that it will not be possible to continue with this pseudo panel past the time period established in advance.

iii. Recruitment methods

To begin with, we contact individuals in a traditional way via a telephone survey. On the basis of specific quotas, we recruit a sample whom we ask to complete an online questionnaire. The use of the internet as the mode of questionnaire administration is, for the target in question, not a major problem. More than 95% of the individuals forming the target are internet users. The distortion associated with the use of the internet can therefore be considered negligible.

The main purpose of the questionnaire is to collect the audience figures for print publications. However, we use this contact for other purposes. At the end of this questionnaire, for example, we introduce the passive audience measurement tool and encourage the respondents to take part in it.

Only mentioning passive measurement at the very end of the press questionnaire is a way of optimising the response rate of the main questionnaire. We were anticipating a significant refusal rate for monitoring using passive measurement. From a financial perspective and in terms of the representativeness of the sample, we could not question only those individuals who would agree a priori to the passive measurement.

At this stage of the survey, therefore, we present arguments for the individual's continuing participation in the full measuring tool. We naturally emphasise the total confidentiality of their answers and of the information passively collected. We also provide the interviewees with a telephone number for obtaining further information about the process, and via which they can also, if they so wish, end their participation in this part of the survey.

Finally, we also use specially trained interviewers to help out any interviewees who may struggle to install the various technical tools that we suggest they use.

iv. Implementing passive measurement

The implementation of the passive measurement system itself is done in several successive phases.

- We start by collecting from the interviewees, at the end of the audience questionnaire, their habits and the dates on which they last accessed the websites participating in the survey. This forms an initial declarative audience base that we will use later.

- We then question the interviewees about the various internet access devices that they possess or have access to. For detailed data, we naturally distinguish the different types of machine (desktop computer, laptop computer, tablet, phone), the type of ownership (personal or professional), the relationship between the machine and the interviewee (whether it is a shared machine or not), and the type of usage (whether the interviewee is the main user or not). This exhaustive report will enable us relativise the audience data later on, and, in particular, to filter out some audience figures measured on devices of which the respondent is not the main user.
- At this point in the survey, we validate the questionnaire, so that we have a good sized sample, even in the event of refusal to participate in the passive measurement. The passive measurement itself is implemented at this stage with the installation of the first cookie. This cookie is placed on the device used to complete the main questionnaire. Of course, we comply with the regulation on the use of cookies which obliges us to inform the respondent about the purposes, methods and duration of the cookie's action. Nevertheless, this same regulation allows us to proceed on an opt-out basis. This means that our duty of information is accompanied by permission to place the cookie when the respondent continues the survey on line.
The ambiguity in the placing of this first cookie also enables us to optimise the rate of installation on at least one device among our interviewees.
- We then invite the interviewees to participate in the rest of the survey by asking them to install a cookie or similar technology on their other devices (an app for mobile devices). We do this by sending an email to the interviewees containing a link for each device that they have declared. The interviewee then has to activate each of the devices in question by clicking on the relevant link for that device.

b. Audience generation tool

i. An incomplete measurement

At this point in the analysis, we already know that we have a tapered panel due to the way in which the sample is recruited. We therefore need to limit the risk of losing the measurement of these individuals during the survey.

To do this, we implemented an activity monitoring tool for the panellists. This monitoring system allowed us to compare, over a month, the list of interviewees who agreed to passive measurement with the list of contacts received, measured on a daily basis. Once identified, these interviewees were integrated into a follow-up process conducted by telephone. In order to ensure its success, we specially trained researchers chosen for their computer skills. The training we gave them involved teaching them to manage the re-installation of a cookie or an app over the telephone with an interviewee.

Another reason for missing data is obviously the reluctance of some interviewees to participate in passive measurement. As we saw above, a cookie is installed quasi-systematically on the terminal used by the interviewee to complete the Audience questionnaire. As such, the rate of installation on at least one terminal is high (85%). On the other hand, the proportion of interviewees who installed the cookie or application on all their devices remains low (13% of the total sample).

Of the criteria that have a positive influence on the installation rate, the most distinguishing criterion is, paradoxically, age. Non-working individuals aged over 65 years presented the highest rates of installation on all their devices (30%). However, this result can be moderated by highlighting the fact that these individuals were also those who used the smallest number of devices, particularly due to the absence of devices used for professional purposes.

Here is the table showing installation rates per device

	Devices possessed	Installed devices	Installation rate
Total devices	45,806	14,679	32%
Computers	22,635	10,077	45%
Tablets	9,546	2,186	23%
Phones	13,625	2,416	18%

Lastly, we must also remember that the measurement, like all measurements that are not systematically linked to the individual, is conducted on the basis of the device, not the individual. It is therefore possible that some of the connections measured come from the device used by a third party. In order to evaluate the potential extent of this phenomenon, we ask the interviewee, when providing a list of devices, to specify for each one whether they are the sole user or if the terminal is shared. We then ask them if they are the main user of the terminal. The table below shows the rates of shared usage and the sharing rate for devices out of all the devices possessed, whether the interviewees agreed to participate in passive measurement or not.

	Desktop computer		Laptop		Tablet		Mobile phone		Total devices
	person.	profess.	person.	profess.	person.	profess.	person.	profess.	
Total devices	7,077	5,615	9,177	4,286	8,711	1,004	9,872	4,564	50,306
Sole user	27%	71%	52%	88%	33%	70%	91%	96%	62%
Shared device	73%	29%	48%	12%	67%	30%	9%	4%	38%
incl. main user	41%	19%	32%	9%	35%	23%	7%	4%	23%
TOTAL sole + main user	68%	90%	84%	97%	68%	93%	98%	100%	85%

This table shows that

- The more mobile a device is, the less likely it is to be shared
- A professional device is less likely to be shared than a personal device
- Only 15% of the devices declared by the interviewees are machines of which they are not the main users.

ii. The principles of injecting missing data

The nature of the data recorded presents certain limitations which were taken into account in the data cleaning phase:

- The measurement is based on the device and not on the individual. With the exception of the smartphone, the terminal is potentially used by several people. Not all contacts are necessarily made by the individual who completed the Audience questionnaire. We therefore defined a cleaning rule which involved removing the visits made from any terminal of which the Audience questionnaire respondent has not stated that they are the main user. 0.6% of visits were therefore removed.
- Some individuals may use several web browsers to access the internet. However, the cookie installed only corresponds to the browser used to complete the Audience questionnaire. Any contacts made via other browsers will not be counted
- The deletion of cookies resulting in discontinuation of measurement should not be disregarded.

Once the cleaning process has been performed, we apply the injection model for the missing data resulting from the non-installation or deletion, intentional or not, of the cookies. However, in order to avoid having to inject considerable quantities of data, we decided to limit the imputation to the period of the month during which each interviewee completed the press audience questionnaire. In this way, we ensure both consistency between press readership and website visits in terms of time scale and preserve the seasonality of website visits.

The imputation model used is the stratified sequential hot deck method. Imputation is performed per title, per device (computer, tablet, smartphone) and per month. All interviewees who participated in passive measurement are potentially donors regardless of their measurement start date (date of cookie or application installation). The data to be imputed is located in the observation: individual x day. In order to take into account correlations between the variables to be imputed and the available variables, the first step is to define the strata within which the imputation is performed. The stratification variables are correlated with the non-response. The sorting variables within the strata are correlated with the variable to be imputed. The variables were selected separately for each of the two spheres of the survey.

In addition to socio-demographic variables, the main stratification and sorting variables are the declarative data on use of the sites in question.

After application of coherence rules, the rate of missing values was estimated at 55%.

iii. Calculation of audiences

The data collected in this way comprise connections between an "individual x device" and a page of a website or application for a given brand. Each of these connections is time-stamped, enabling us to create a log of page views.

Using these pages we establish the visits, unique visitors and times on page using the rules applied for site centric measurement by the certification organisations such as ABCs:

- A unique visitor means all deduplicated daily visits by one individual x device.
- A visit means all the pages viewed during a single, uninterrupted succession of page views on one site or application during a single session.
- A session of visits means all the visits made in an uninterrupted manner. A duration of over 15 minutes between two pages constitutes an interruption of the session
- Time on page is determined by the time spent between the opening of two pages. It is therefore not possible to calculate the time on page for the last page of a session.

From the numbers of unique visitors determined in this way, we obtain the daily basis on which the injection is performed. After this, we establish the monthly audience values per site and per device.

III- THE PRINCIPAL RISKS TO OVERCOME TO ACHIEVE GOOD-QUALITY MEASUREMENT

On the basis of the methodological principles set out above, it is clear that the implementation involves many risks of which you need to be aware, in order to avoid damaging the measurement.

a. Publisher-related risks

i. Difficulties implementing tags and SDKs

The primary difficulty for publishers is the correct installation of tags on websites and SDKs on applications. Although tags do not present a major technical problem, they must nevertheless be positioned on the site such that they cover the whole site but, within the context of press groups, do not attract non-legitimate audiences, such as those from "channels" that do not form part of the principal brand.

For this last point, we collected the URL of the page viewed at the same time as the visitor connection logs. Using this we were able to constantly validate the scope of the measurement.

The installation of the SDK within apps poses other problems. The first issue is that the app has to be republished after the installation of the SDK. This also means that the new version has to be accepted by the stores, including the Apple one.

Given that the sites measured are all commercial sites and that our technology partner for this operation is also an advertising server, there was no publication refusal from Apple. Nevertheless, this point could constitute an obstacle.

ii. Tag removal and app republication issues

We were also confronted with issues relating to the IT services responsible for publishing sites. Some of them found that the use of tags on their sites slowed page loading times. They therefore removed, without warning the tag that we were using for the survey. Luckily, the tag monitoring system enabled us to limit the effects of this type of situation within a reasonable time-scale. However, this again is a major consideration that could render the measurement of a site completely ineffective.

Although the absence of a tag can be easily detected, the same is not true for SDKs. If the publisher of an app wants to update it, they will not necessarily reintegrate the SDK that we have installed. The reason for such an omission is usually the fact that the publishing teams are not heavily involved in audience measurement, which is usually the responsibility of the advertising companies or the marketing department. In these circumstances it is easy to omit the SDK.

And while it is easy to omit it, detecting the omission is particularly difficult. Once an app is downloaded on to the user's phone or tablet, it remains active with full functionality (including the SDK) until it is updated. So, even if the publisher forgets the SDK in a new version, we will continue to receive connections due to the existence of old versions of the app still in use. We therefore need to be particularly aware of situations in which contacts gradually fall below their usual level.

These points, as incidental as they may seem, are essential to the proper measurement of the digital audience. Every day for which a tag or SDK is absent will generate new missing data in the connection database.

b. Panellist-related risks: ad blockers and cookie deletion

The main reason for the cookies not being installed by the interviewees is existence of an ad blocker on the device used to complete the Audience questionnaire. Ad blockers were present on 15 to 20% of the devices on which we tried to place a cookie. At the same time, these devices correspond to over 60% of the causes of failure to place a cookie on devices, the other reasons mainly being refusals to participate.

This point is all the more damaging to the survey given that not only does it prevent passive measurement, it may also, in certain cases, block the questionnaire completely.


In order to avoid this situation, at the start of the questionnaire we introduced a system for detecting the presence of an ad blocker. If it detects one, the interviewee is offered an alternative version of the questionnaire which does not attempt to install the passive measurement system. This avoids losing otherwise complete questionnaires and limits the risk of collecting erroneous information during the passive measurement phase.

In addition to installation failure issues, we also had to manage situations in which cookies were removed. These cases are much more difficult to identify. The deletion of a cookie resembles a non-connection. It is, of course, possible to conduct a telephone follow-up campaign to invite the panellist to reinstall the cookies. However, if the panellist is simply not very active on the internet, this follow-up may, at the very least, be considered unwelcome and, in the worst case, may result in the panellist deciding to discontinue their participation.


We therefore decided to be very cautious during this follow-up phase and to limit it to individuals from whom no connection has been received for at least one month.

The deletion of cookies, where it is an actual technical issue, is actually quite rare. We were able to evaluate it by comparing two normally correlated variables: the number of panellists recruited and the number of connections detected per day. If the two did not correlate, we would have evidence of significant deletion of cookies.

The graph below shows, on a daily basis, the number of connections measured (red curve) compared with the total number of panellists recruited to date (blue curve).



You can see the clear overall correlation of the two curves but also the significant irregularity of the connections curve. On the other hand, if we remove from the graph weekends and public holidays (when there are no connections from computers in the workplace), and days on which one-off events take place (attacks, sporting events, etc.), we see an almost perfect alignment of the two curves. This shows that the underlying volume of connections to press sites is totally regular and proportional to the volume of panellists monitored.



IV- COMPARATIVE ANALYSIS OF AUDIENCE FIGURES PRODUCED VIA DECLARATIVE AND PASSIVE MEASUREMENTS

Lastly, since our survey is based on press site connection data that is collected in both a declarative and a passive manner, we analysed the possible relationship between these two modes of measurement.

This analysis is particularly interesting given the substantial volume of passive data collected. We have 3600 panellists active on press sites who have generated over 27,000 visits and a total of more than one million page views.

This analysis was, of course, performed on the raw data before any injection of missing data in order to use the purest information possible. The declarative data were collected for each of the sites taking part in the survey, on the three internet access devices (computer, tablet and mobile phone), based on two indicators:

- Site connection habits on the following basis
 - o Daily
 - o At least once a week
 - o Two to three times a month
 - o Once a month
 - o Less often
 - o (Never)
- The date of the last site connection on the following basis
 - o Yesterday
 - o Less than 8 days ago
 - o Between 8 and 15 days ago
 - o Between 15 and 30 days ago
 - o More than 30 days ago
 - o (Never)

As regards the data collected using passive measurement, we have, for the same individuals, the number of connections detected on each type of device for each press site.

By comparing these two sets of data, we learned a lot about the relationship between declarative and passive data. The first thing we learned concerns the number of connections to non-declared sites. Over all the sites measured, we observed that 36% of visits were made by individuals who declared that they never visit the site or application in question. Many visits are made without remembering the site visited.

On the other hand, if we look at the individuals who declared a habit, there is a high concentration of visits measured passively for the highest habit levels:

	Daily	At least once a week	2-3 times a month	Once a month	Less often
Spread of visits	66%	16%	8%	4%	6%

If we look only at the concentration in the first habit level, we see that it is higher on mobile applications, whereas the other connection types are similar for this criterion. It is likely that the strong identification of visits via mobile app is due to the very personal nature of the telephone and the conscious choice to install an app.

	Share of visits on "daily" habit
Internet Site	64%
Mobile Site	64%
Mobile App	81%
Tablet App	61%
Tablet Site	63%

There is therefore a link between declared habit and the number of visits recorded. What about the criteria of "recency"? Here again we see a strong link, with 84% of the visits declared by the panellists corresponding to visits made during the last week.

	Recency of visit				
	Yesterday	Less than 8 days ago	Less than 15 days ago	Less than a month ago	More than a month ago
Spread of visits	67%	17%	6%	4%	6%

Based on these results, it seems likely that comparing the two variables of habit and recency would produce a similar result. And this is indeed the case. We even see a concentration of 73% of visits within the first level for at least one of the two variables and 80% within the area comprising the first two levels for both variables.

		RECENCY				
		Yesterday	Less than 8 days ago	Less than 15 days ago	Less than a month ago	More than a month ago
HABITS	Daily	61%	5%	0%	0%	1%
	Once a week	5%	9%	1%	0%	0%
	2-3 times a month	1%	3%	2%	1%	1%
	Once a month	0%	1%	1%	2%	1%
	Less often	0%	0%	1%	1%	3%

This overall analysis is instructive, but we wanted to go further by analysing the results based on the same principle but focussing on the main categories of press.

We then see that not all press categories are equal in terms of the value of declarative data. News-related sites naturally show a strong link between the declarative and passive measurements.

	Share of visits on undeclared sites	Share of declared visits according to frequency				
		Daily	At least once a week	2-3 times a month	Once a month	Less often
National newspapers	25%	70%	15%	6%	5%	5%
Regional newspapers	34%	67%	18%	7%	4%	5%
News magazines	49%	39%	25%	16%	11%	9%
Economics magazines	64%	34%	20%	27%	10%	9%
Entertainment magazines	81%	25%	17%	29%	14%	16%
Women's magazines	82%	22%	29%	17%	21%	11%

CONCLUSION

This experiment has enabled us to show the opportunities, risks and benefits of introducing passive measurement of internet connections into a traditional audience survey.

- **Feasibility:** we have demonstrated that
 - o The use of cookies and tags is an operational option for this kind of measurement
 - o The lifespan of a panel formed in this way was sufficient to perform usable measurement
 - o A panel formed from a sample can generate a sufficient quantity of data
 - o Interviewers were able to improve digital data response rates with training on providing assistance with cookie/application installation
 - o We also managed to install a sufficient number of devices to conduct this type of passive measurement properly
- **Risks:** we have discussed the main points to watch out for, including
 - o The logistics for installing cookies and tags, which must be perfect
 - o The issue of the existence of ad blockers on the panellists' devices
 - o Ensuring that injected data is good quality, otherwise the information generated is no longer valid
- **Challenges for the future:** in our opinion, the most important issue to be faced in future is the development of e-privacy rules. The new restrictions to be applied in Europe in Spring 2018 will significantly limit opportunities to compare data and, as such, the capacity to perform this type of measurement properly.
- **The relationship between declarative and passive data:** we demonstrated two key points
 - o Declarative data can only be used to collect some of the press website visits, with passive measurement providing information on additional visits that the panellists did not remember
 - o The data missing from the declarative measurement varied between press categories
 - o On the other hand, the declarative data on habits and recency are a highly predictive indicator of the number of actual connections to the site in question. As such they can be used as relay variables to bridge passive measurement gaps for individuals reluctant to participate in a panel.