Chandra Users' Committee Meeting Preliminary Report

2020 October 15-16

Overview

The CUC continues to be pleased with the observatory performance and management, as well as the dedication of staff to supporting the user community. The committee especially commends the CXC staff for their sustained efforts and high-level performance throughout the most recent, world-wide challenges, including swiftly adapting to teleworking. We also wish to thank Dr Belinda Wilkes for her leadership and commitment to the mission, and to extend a warm welcome to Dr Pat Slane, who recently took over as Director.

Chandra Status Report

From the perspective of program management and science and mission operations, the Chandra X-ray Observatory and the Chandra X-ray Center are performing well. Staff have adapted to the coronavirus pandemic by implementing teleworking, with minimal necessary on-site presence. The CXC has conducted events such as the Cycle 22 peer review and the Chandra science workshop on time-domain science remotely, and have successfully carried out important operations activities while teleworking. Fiscal year 2020 expenditures are in line with budget guidelines, although margins are small and the CXC has at times had to deal with incremental funding. Chandra's science instruments are performing well overall. A hardware anomaly in the HRC was resolved by swapping to the instrument's redundant set of electronics; further checkout is in progress. Maintaining temperature control of the spacecraft continues to make mission planning complex, and additional challenges are anticipated due to the orbit's decreasing perigee in coming years.

The CUC is very pleased to see that spacecraft and instruments work very well, maintaining a stable, high efficiency.

Director's Report

Former CXC Director, Dr B. Wilkes, reported on the Director's Discretionary Time allocation over 2020/2021. Overall requests have increased; DDT now around 1 Ms, and requests use most of the available time. The current system for scheduling requests is labor-intensive, but appears to be working well. Many programs included multi-wavelength campaigns impressively coordinated across observatories.

An overall summary of the DDT targets was shared, including four non-transient targets. Six interesting and exciting objects targeted during DDT were presented to the committee, spanning a range of science from a young star with a protoplanetary disk to symbiotic stars, AGN, pulsar wind nebulae, and X-ray binaries. Betelgeuse was visited in its dim period in early 2020, but after two non-detections, the monitoring was concluded.

The Committee was happy to see the two new books celebrating 20 years of Chandra, "Light from the Void" and "The Chandra X-Ray Observatory" were published at the end of 2019 as scheduled. Dr. Wilkes notes the latter is a very handy reference text.

The CUC is pleased that as in previous cycles, DDT is competitive, a broad range of science topics is covered, and high-impact science results are produced via these observations.

Mission Planning Updates

The goal of Chandra mission planning is scheduling of observations to maximize science return in the presence of observing and engineering constraints. The evolution of the spacecraft's thermal properties poses one of the major difficulties for this goal, as observations cannot be freely scheduled over the full sky area accessible to Chandra, but must be sequenced to respect complex and growing thermal constraints. Management of spacecraft angular momentum is an additional complication of growing concern. Mission planning is maintaining high observing efficiency, meeting observation constraints, and operating a robust program of TOOs and DDTs through a number of mitigation strategies, from improved computational tools for scheduling and more effective TOO and DDT processes to more accurate ways of assessing and limiting observation difficulty for use by observers and the peer review process.

The CUC is very happy to see that the impact of the degradation of the spacecraft insulation on scheduling appears to have been smaller than expected. We commend the CXC on this success and the changes in operation, planning and software that have made this possible. We strongly urge the CXC to continue to put substantial effort into this area.

From the response of the CXC to our recommendation of a "Be Cool" program, we understand that the CXC does not expect thermal issues to have a major impact on scheduling in Cycle 23 and perhaps Cycle 24 as well. We would nonetheless recommend the resource cost of a target be displayed automatically as a target is entered into the program and/or that a tool be provided that could take a list of sources and report their resource costs directly. This would enable users to include the expected resource cost in their choice of target even before they enter them into their proposal.

Proposal Cycle and Future Plans

Dr Prestwich gave an overview of the status of Cycle 22, plans for Cycle 23 and highlights of the past few months. Regarding Cycle 22, the number of proposals was very similar to Cycle 21. A total of 164 proposals were approved, including 15 LPs and one VLP. No negative trends were noted with respect to gender statistics (whereas the number of senior female PIs tends to be lower than their male counterparts, that effect is counterbalanced by a high degree of success amongst junior female PIs). The number of LPs approved is higher than normal due to a larger allocation of time.

Several initiatives were introduced in Cycle 22 to facilitate thermal management of the spacecraft and enable a stable long-term schedule. These include: the implementation of a Resource Cost metric, elimination of observing preferences, changes to TOO response times and follow-ups, and a 4 Ms increase in observing time.

The "Resource Cost" (RC) metric is a point system that accounts for limits on the available observing time at high ecliptic latitude. RC essentially replaces the "Constraint categories" used in previous Cycles, such that targets at high ecliptic latitude will incur significant RC even if they are not constrained. While this system worked well, and yielded no cuts in approved targets, the issue did arise that a panel met their quota before allocating all of their time. In light of the increased challenges with the telescope cooling, CUC had previously recommended considering the creation of a “Be Cool” program, where a score can be assigned to gauge how useful a target will be for cooling the telescope (not necessarily for grading purposes). This concept may be implemented in the future.

CUC had previously endorsed the creation of a working group to look at the best way to handle the expected increase of requests for transient science, and the consideration of a related workshop. Owing to the ongoing public health crisis, this morphed into the development of the online workshop “Chandra Frontiers in Time Domain Science”, which is taking place at the time of writing. A dedicated working group will collect ideas and suggestions from the workshop participants with the specific purpose to make recommendations on how to best handle high-impact, transient CXO science.

Also a result of the ongoing public health crisis, in 2020 the CXC was forced to switch to a fully remote Peer Review for Cycle 22. The schedule was modified to accommodate remote participation: the discussion times were limited to 5 hours per day, and the review lengthened by 3 days to account for shorter discussion times per day. Zoom was used for the meeting software and Slack. Each panel had a facilitator who managed both the meeting software and the panel database software. Responses to a reviewer survey shows that most reviewers felt that the format worked well, with an even split between people who prefer in-person and remote reviews. Considering the ongoing uncertainty, Cycle 23 will also take place remotely. Additionally, the Dual-Anonymous Peer Review (DAPR) process mandated by NASA Astrophysics will kick off in Cycle 23. Trained “levelers” will observe the panels' discussion and intervene for the sole purpose of refocusing the conversation on science, and stop it in case of egregious violations.

The CUC appreciates that 2020 brought a unique set of challenges, including but not limited to switching to a remote review etc. The committee commend the CXC staff for their sustained efforts and performance throughout an extraordinarily challenging year.

Recommendations:

1. *The committee renews our recommendation that the Director keeps exploring the possibility of Joint Observing programs with ALMA.*
2. *We recommend the addition of "pull-down", readily visible Resource-Cost on a target by target basis to be implemented for the next proposal Cycle.*

Calibration: Goals, Priorities and Plans

The CUC received a comprehensive presentation regarding the multiple activities performed by the CXC. Several instrument performance parameters are time-dependent and require moderately frequent monitoring, such as detector gain, quantum efficiency, and composition of the contamination known to be building up on the detector surfaces.

ACIS: Cold ECS data has always been the fiducial point of ACIS calibration. Due to the fading of the ECS (2.7 year half-life) and the lower fraction of cold ECS data, the calibration team estimates that the ECS can only be used for calibration for one or two more years. The calibration team carried-out a proof-of-concept set of observations on I3 to determine if Cas A can be used to calibrate the ACIS gain in the future. The results look promising. All basic properties of the ACIS CCDs (gain, QE, spectral resolution) depend on the focal plane temperature with the FI chips are much more sensitive to temperature changes than the BI chips. The calibration team has approved the scheduling of S3-only observations up to -109C. However, it has requested that the present limit of -112C be maintained for FI chips. The calibration team has begun generating a set of calibration products for a grid of temperatures. The current CALDB version of the contamination model slightly underestimates the depth of the contaminant on the ACIS filters at the present time. An updated contamination model will be released by Dec. 15, 2020.

HRC: The calibration team continues to monitor the gain and QE of the HRC detectors through annual observations of HZ43. Annual updates to the HRC-I and HRC-S QE gain and QE continue to be released. Correcting the raw HZ43 count rates with the appropriate QE files produces a very small scatter in the fluxes of HZ43 of only 2%.

HETG: The effects of warm ACIS data have little impact on HETG/ACIS-S data since the photon energies are computed by the dispersion relation. Therefore, the calibration team has approved to carry out HETG/ACIS-S observations up to -109C.

HRMA: AR Lac is observed with the HRC-I twice per year to monitor the PSF. The last two observations showed a slight increase in the 90% encircled energy (by 0.04"). The calibration team will continue monitoring the PSF, and if this trend persists, a plate focus test may be needed.

Recommendations:

1. *The CUC continues to strongly recommend that the current calibration efforts are maintained at the same level of commitment. The long-term trends seen in many of the detector parameters require frequent monitoring to ensure the continued success of the mission. The CUC, therefore, again asks the CXC to continue providing regular updates with the latest data on the evolution of these properties at future CUC meetings.*
2. *The CUC again recommends that the CXC calibration team continue to play an active role in the IACHEC collaboration in order to maintain efforts to characterize and understand the cross-calibration uncertainties between the various X-ray missions currently active (e.g. XMM-Newton, NuSTAR, Swift, NICER).*

Chandra Source Catalog

The second major release of the Chandra Source Catalog (CSC 2.0) includes ACIS and HRC-I imaging observations released publicly through the end of 2014. The sensitivity is maximized by co-adding observations in each region of the sky. CSC 2.0 includes 317,167 unique X-ray sources and ~928K individual observation detections (~1.42M entries including photometric upper limits). The catalog consists of tabulated source properties and upper limits, and science-ready FITS format data products for further analysis (both full-field and source/detection-based). The CSC 2.0 data are served to the community via multiple interfaces, and are actively used in scientific research papers, covering a wide range of subjects from stars to cosmology.

The CXC is developing a minor update (CSC 2.0.1) to correct errors in several source-level inter-observation variability properties. This update will be released to the community by the end of the year. Work is also nearing completion on cross-matching the CSC 2.0 source positions with the SDSS, Gaia, ALLWISE, and Pan-STARRS catalogs: ~150,000 matches were found and will be made available to the community. Some of these sources will be observed as part of the SDSS-V spectroscopic time-domain survey. The CXC will support a local public repository for these optical spectra.

After completion of the ongoing fifth reprocessing of the entire Chandra observation set, expected to complete in late 2021, the CXC plans to start processing of CSC release 2.1. This release will add observations released publicly from 2015 through 2021 to the catalog. CSC 2.1 will tie source astrometry to the Gaia reference frame and resolve some issues found during CSC 2.0 production by using more robust algorithms. The expected processing time is ~1 year to completion. Prior to the final release, new data will incrementally be made available via the CSCview Current Database view as catalog production proceeds. Processing of observation fields that do not overlap existing stacks will be prioritized.

The CUC was pleased to hear that, regarding the eventual processing of CSC2.1, plans are in place to prioritise processing of 'new sky' that does not have significant overlap with the sky area already processed for CSC2.0, in line with its previous recommendations.

Recommendations:

1. *As in previous reports, the CUC once again strongly recommends a strategy in which reprocessing new regions of the sky for CSC2.1 takes the highest priority (following the completion of REPRO-V), and efforts are focused on catching up on the processing of publicly available data. The CUC was naturally disappointed to hear of delays in this process, although it is understood that this is largely due to the effects of COVID, which could not be foreseen. However, given the well-documented history of the CSC, the CUC was also concerned that new features/updates appeared to be being discussed for implementation within CSC2.1, which may carry the risk of further delays. The CUC therefore also wishes to reaffirm its previous recommendation that priority be given to processing the new data with the existing CSC2.0 framework, and updates to this framework only be made if they are unavoidable (e.g. compatibility issues with updated hardware) and/or address genuinely critical issues.*
2. *The CUC also recommends strongly that a detailed plan be put in place for regular future releases of versions of the CSC to keep them up to date. Such a plan should include a specific frequency (e.g., monthly or twice yearly) for when new releases can be expected.*

SAO membership in SDSS-V

Dr. Paul Green reported that SAO is signing up as a full institutional member of SDSS-V (with half of the dues being contributed by the CXC). SDSS-V is a 5 year, all-sky multi-epoch optical/IR spectroscopic survey that will provide spectra for millions of stars (Milky Way Mapper), about half a million quasars and X-ray source counterparts (Black Hole Mapper) and IFU spectroscopy of the Milky Way and local galaxies (Local Volume Mapper). CXC scientists at SAO and MIT now have full access to SDSS-V data and collaborations, and can include postdocs, grad students and undergrads in their SDSS-V-related research.

SDSS-V will target some 40,000 X-ray sources identified from the Chandra Source Catalog. CXC will publically serve updates containing those targeted spectra and pipeline data characterizations (e.g., star, stellar spectral type, galaxy, quasar, redshift) to the community on a regular schedule. This product will be complementary to but served separately from the Chandra Source Catalog.

The CUC is pleased to learn about this development and the value added to the Chandra Source Catalog provided by the IR/optical counterparts to ~40,000 X-ray sources.

CIAO updates

Dr. Jonathan McDowell gave an update on the CIAO software. Downloads of CIAO have been stronger than ever and the mix of operating systems is still evenly distributed among the users. A single Mac-OS and Linux release was implemented for CIAO 4.12 and an option of installing CIAO using conda was also introduced. The conda install will probably be the default installation for future CIAO releases. In fact, SDS and DS plan to release CIAO 4.13 in December 2020 with a python 3.8 default version for conda installation, but ciao-install will still be available.

Routine updates to the documentation have been made, including infrastructure improvements to better support Jupyter notebooks and to unify the Sherpa documentation. Chips (which will be retired in the next CIAO release) has been removed from the current documentation and matplotlib is now fully implemented in CIAO and Sherpa analysis threads. A significant enhancement of the DAX menu in ds9 now provides a powerful yet simple GUI interface to CIAO tools. A replacement for the old Prism tool with a new Prism within ds9 is currently under development and the upcoming release of CIAO 4.13 will include an experimental version of Prism in ds9.

The CIAO team continues to actively support direct outreach to the community with CIAO workshops. A successful workshop (in person) was organized at the AAS235 meeting in January 2020 and a fully virtual one is planned for the upcoming AAS 237 meeting in January 2021. A remote workshop is also planned in Egypt, with approximately 60 participants from various Middle East countries. Social media outreach is also steadily increasing.

The Committee thanks the CIAO Team for their commitment to development and further improvement of the software and for continuing to address Helpdesk tickets in a timely and effective fashion. The Committee also commends the team for their sustained efforts to support direct outreach to the community and adapt to current travel limitations caused by the ongoing pandemic by organizing fully remote CIAO workshops. In particular, the CUC is very pleased to hear that a virtual workshop will happen in Egypt and thus reach the Middle East community.

Chandra Exclusive Data Use Policies

The CXC has been requested by NASA HQ to review the current policy for exclusive-use periods (EUPs) for Chandra data. This is in view of the point that this period of 12 months for data from Chandra General Observer (GO) proposals, but only 6 months for data from HST. Specifically, CXC was asked to address the need to maintain 12 month proprietary rights while HST recently reduced it from 12 to 6 months only.

General comments: NASA Missions appear to adopt a range in EUPs. Explorer-class missions such as Swift and NICER have no EUP; all data become publicly available immediately after entering the respective archives. The same is true of Fermi, which is almost exclusively a survey and monitoring mission. Chandra, NuSTAR, and JWST have 12 month EUPs, as did HST for the first 25 years of the mission, after which it was changed to 6 months. We wish to stress that this happens at the same time when the JWST, which is widely seen as the natural "successor" to HST, was set to fly. Unfortunately, a successor mission has yet to be identified when it comes to Chandra, whose spatial resolution is X-rays is vastly superior to XMM, its ESA counterpart.

The EUP for Chandra observations is different for different programs. For GO, GTO, and LP programs, the period is 12 months. For DDT programs, the period is 0-3 months, determined by request and the decision of the Director. VLP and XVP programs carry no exclusive use period.

The Chandra policies for EUPs have been discussed in multiple previous meetings with the CUC, and in the 2016 Senior Review. There has been broad support for the current policy, with acknowledgment that this should receive future reviews. Most recently, there have been suggestions that "high-profile" targets (generally transients) should have shorter EUPs. From the minutes of the 2019 CUC Meeting:

"In some cases there are significant differences in proprietary times between the major observatories (in particular HST and Chandra, but also the VLA and ESO) for (typically transient) objects of great community interest, such as GW170817, so as to enable rapid, multi-wavelength follow-ups. We recommend that the Chandra Director approach the other observatories to see if a common set of rules could be adopted."

The CXC is currently hosting a Workshop on "Chandra Frontiers in Time-Domain Science" and will convene a group to discuss issues and suggestions raised in the Workshop to address this and related ideas. There are multiple factors to consider in evaluating the appropriateness and desirability of EUPs for Chandra (and other) data, including impacts on students, young faculty, and faculty at small teaching colleges.

A study of Chandra-based peer reviewed publications indicates that the time delay between completion of observing programs and the resulting publication(s) is roughly two years for both programs with standard EUPs and those with none. This is consistent with results pertaining to the HST publications as well. In short, removing EUPs does not lead to faster publications. At the same time, faster publications at the expense of incomplete analysis due to not using full datasets is not a desirable result.

Because long Chandra observations are necessarily split into smaller segments, sometimes distributed over the course of an entire year, complete removal of an EUP for such programs would offer significant temptation for other groups to publish intermediate results before the PI team obtains the full data upon which the accepted program was proposed. As above, this is generally not considered a desirable result.

In summary, the CUC did not reach full consensus on this topic. Overall, a reduction from 12 months to 6 months for the exclusive use period for all Chandra observations was considered to provide no strong benefit to the larger community. However, there was some feeling that a reduction (or even elimination) of the exclusive use period might be beneficial for some time-domain observations where immediate follow-up observations at other wavelengths could enhance the science return.

Arguments against reducing the proprietary time period fall under the following major categories:

- *Longer proprietary time periods might enable better science.*
- *Longer proprietary time periods protect students from competition.*
- *Longer proprietary time periods protect faculty with heavier teaching duties from being scooped.*

The arguments in favour of reducing the proprietary time period can be summarized as below:

- *A step towards making Astronomy more inclusive: "open data" effectively means sharing resources with the entire scientific community in real time, including scientists at institutions that historically do not get regular access/support through GO programs.*
- *This could also enable radically new science, e.g. the mining data in real time enables the identification of transients that may serendipitously appear in the data stream (without the need to implement new policies).*

We end by noting that the Committee member views on proprietary rights might not be necessarily representative of the entire community, even within the time-domain community. We recommend that a dedicated working group is created to further explore this issue.

Somewhat separate to the issue of proprietary rights, the CUC recommends exploring options for mining Chandra data as new observations are acquired, so as to readily identify new X-ray transients in real time. The CUC understands that this presents numerous challenges, including, but not limited to (1) the feasibility of data mining from a data-analysis and delivery perspective; (2) the implementation of new policies to allow proprietary data to be mined that do not interfere with the interests and rights of the PI (that is, in the presence of proprietary rights).

HRC Anomaly and Current Status

The CUC received two detailed reports from Dr. S. Hurley and Dr. D. Patnaude regarding the recent anomaly seen in the HRC, covering the response of the CXC to the event and the current status of the HRC detector.

On 24th August 2020 an anomalous power configuration was seen and identified as an issue with the HRC. Although the issue was confirmed to be limited to the HRC system, this currently provides the primary anti-coincidence shields that help to ensure the safety of all of the *Chandra* detectors (including ACIS) against periods of high radiation. The need to power down the HRC to address the anomaly therefore also resulted in a period of time where it was considered unsafe to operate the ACIS detectors, and in turn a temporary cessation of science operations. In total, all of the relevant systems were powered down within an ~hour of the anomaly initially being discovered.

Subsequent investigation into the cause of the anomaly after the power-down revealed that the issue was in the HRC A-side electronics box. In order to test whether this was a single event issue that could be cleared with a reboot, the A-side electronics were powered up again. Ultimately, though, the anomaly did reoccur, and the HRC was once again powered down to facilitate a switch to running the HRC with the B-side electronics. This switch was successfully made on 31st August, but does now mean there is no further redundancy regarding the HRC electronics systems. All systems were gradually powered up again over the next few weeks, with performance closely monitored. Both the B-side HRC electronics and the anti-coincidence shields were seen to be performing as expected, and so science observations with the ACIS detector were re-commenced on the 12th September. This is the largest period of down-time the mission has experienced to date.

Efforts are now ongoing to re-commence science operations with the HRC, as ~1.1 Ms of HRC observing time was awarded in *Chandra* cycle 22, some of which is time sensitive. The voltage for the HRC-I has already been manually ramped up in order to conduct a real time observation of AR Lac (a known calibration source). These data have been independently analyzed by members of the SOT, instrument, and calibration teams, and the HRC-I appears to be functioning nominally again. A similar activity is planned for the HRC-S in the coming weeks. Additional checkout observations, one of Cas A with the HRC-I, and one of Capella with the HRC-S/LETG, are also planned. Once it is confirmed that the HRC is functioning nominally, it will be returned to normal science operations. An initial estimate for this timeline is ~mid November, assuming further re-calibration efforts are not required.

The committee would like to strongly commend the HRC, ACIS and flight teams at the CXC for their rapid and decisive response to this anomaly, and for their ongoing efforts to fully recover Chandra science operations, particularly under the currently challenging circumstances relating to the coronavirus pandemic.

Recommendations:

In light of the recovery efforts so far, and those still planned, the CUC did not have many additional detailed recommendations. However, a couple of broader points were noted:

1. *The CUC recommends that efforts to fully recover science operations of the HRC are continued in as timely a manner as possible (without introducing unacceptable risk to the detectors), particularly given that some of the science observations approved for cycle 22 are time sensitive.*
2. *The CUC also recommends that efforts be continued to establish contingency plans for protecting the ACIS detectors in the instance that an issue with the HRC B-side electronics arises in the future, and the HRC anti-coincidence shields cannot be run with the degraded A-side electronics.*

NASA Hubble Fellowship Program

Dr. Paul Green reported on the NASA Hubble Fellowship Program (NHFP). The program is sponsored by NASA, administered by the Space Telescope Science Institute, and scientific leadership is shared among three leads representing the three science "flavors": Hubble (Andy Fruchter, STScI), Sagan (Dawn Gelino, IPAC) and Einstein (Paul Green, CSC). Since 2018 the program covers a wide range of astrophysics of interest to NASA's missions and unites three original categories of fellowships. In January of 2020 twenty four fellows were selected out of 380 applicants during a review that involved 50 reviewers, 7 panels and 8 chairs. One recipient eventually declined the fellowship, bringing the overall number of fellows to twenty three (including 8 Einstein, 7 Hubble and 8 Sagan). This is comparable to the numbers from the previous year, when twenty four fellows were selected out of 383 applicants (7 Einstein, 11 Hubble and 6 Sagan). The Symposium for the 2020 cohort of fellows was held remotely and included both science and non-science sessions. The latter were informed by interests and preferences of the fellows and included sessions on NHFP Against Anti-Black Racism, discussion of benefits and policies, and an open mic session.

The recent policy changes to the Program now allow (a) application up to four years post-PhD for extenuating personal circumstances (like childcare) and (b) starting with 2021 cohort, the Program will require that all NHFP hosts must offer employee status to fellows (as opposed to stipendiary status). This option will also be extended to existing fellows from earlier years. Other policy changes presently under consideration include: support for mentoring activities (e.g., mentee graduate student travel to observatories) and for mentoring of students of color, support for visits by collaborators, expansion of host locations to include the NASA research centers (e.g., Goddard, Marshall, Ames).

The due date for applications for the next cohort of 2021 NHFP fellows is November 4, 2020, with a November 12 deadline for the letters of reference. The review of applications will be fully remote and is scheduled for January 2021. The recruiting of reviewers for the new cycle is nearly complete, the software is in place and fully functional, and the NHFP program leads anticipate a smooth review cycle. The Program continues to distribute fellows to institutes throughout the US with rules on the number of fellows resident at each institution (specifically, no more than 2 new fellows per institution in any given year and no more than 5 fellows total at any given host). The review of the Program is scheduled for the spring 2021.

The CUC is pleased to see that the NHFP continues to provide a supportive community through annual Symposiums and is actively seeking ways to improve the fellowship in terms of its benefits and policies.

We commend the NHFP scientific leads for actively involving fellows in the discussion about how to be more family-friendly and inclusive.

Recommendations:

1. *The CUC reiterates a recommendation made last year: "To encourage people to apply to the program, statistical information on the number of applicants and those selected should be posted on the web."*
2. *In addition, we recommend that the Program keeps records about the representation of fellows from different (over- and underrepresented) groups in proportion to the number of applications from their group, and to use them as guidelines to ensure equity. Because the NHFP reviews are not double-anonymous, it would be beneficial to understand the impact of the prestige of institutions from which fellows originate (their PhD awarding and/or first postdoc institution) and the impact of the placement of their letter writers (institution and geographic location) on their chances to be selected. It would also be of interest to understand the distribution of institutions where fellows choose to go.*