LIGO

LIGO LABORATORY
MS 18-34
PASADENA CA 91125

TEL: 626.395.2129
FAX: 626.304.9834

MEMORANDUM

DATE: January 18, 2016

TO: Fulvio Ricci and Gabriela González
FROM: Stan Whitcomb, on behalf of the Detection Committee
SUBJECT: Report of the Detection Committee
Refer to: LIGO-T1600014-v1

The Detection Committee was convened by the LSC and Virgo spokespersons and consists of

Bruce Allen	Frédérique Marion, co-chair	B.S. Sathyaprakash
Enrico Calloni	Nergis Mavalvala	Peter Saulson
Fabien Cavalier	Brian O'Reilly	Jo van den Brand
Vicky Kalogera	Keith Riles	Andrea Viceré
Mike Landry	Sheila Rowan	Stan Whitcomb, co-chair

The charge to the committee is reproduced in the Appendix to this report. Since shortly after GW150914 was recognized, the Detection Committee has been working in accord with the Detection Process described in LIGO M1500042/Virgo VIR-0103A-15 (The Process for Claiming a First Detection). Early activities focused on evaluating the state of the instruments and the data quality. The Detection Committee has evaluated these via review of the DetChar Detection Checklist, through interactions with the Rapid Response Teams from the sites, and through interaction with the detector design and commissioning teams. After the data analysis groups consolidated their findings, we interacted with them and with the Paper Coordinating Team to evaluate the detection case and how it is presented. Finally, we have reviewed the available drafts of the companion papers asking whether the information in them is consistent with that in the detection paper, and whether they reinforce our main themes.

In the report below, we respond to the elements of the charge.

On behalf of the full Detection Committee, I would like to thank everyone in the two Collaborations. The hard work of the instrument teams, the operational staff at Hanford and Livingston, the detector characterization group, the data management teams, the analysis groups, and especially the Paper Coordinating Team has made this culmination to a long campaign possible. Everyone has been very responsive to our requests for clarification and further studies. We could not have completed our work without their help.

Summary of Findings

On the primary question in the charge (do we believe that the case is sufficiently strong to make a detection claim?), the answer from the Detection Committee is a resounding and unanimous "yes".

We are convinced that the detectors were operating normally at the time of GW150914. Even though the event occurred before the formal start of O1, both detectors had already achieved stable operation with better-than-expected sensitivity. The configurations of the detectors are well documented and their noise behaviors are well characterized. There is nothing in the hardware, the environment or the operational records that would cause any doubt regarding the validity of what was measured.

The analysis pipelines have performed as planned, even with the challenges that arise from this event occurring before any substantial block of background data was collected. The analyses have had to deal with creating an unbiased background even though the participants knew that there was a candidate event under consideration. The search groups have performed well, and the fact that multiple independent searches returned this candidate with high significance gives us great confidence in its validity.

At the time this result is announced, readers of our paper will know that we have additional data that we are not reporting on, and it will be natural for them to ask what we know about candidates in the rest of the run. We believe that a certain measure of openness about this topic will strengthen our credibility. We recommend adding a sentence to the concluding section of the detection paper along the following lines:

> "Data from the remainder of the O1 run are currently being analyzed. At least one additional candidate has been identified, but the careful checks needed to validate this event have not been completed. It will be published if it passes these checks."

During the time that the paper is under review, we should continue to perform the data quality checks to validate this event. The suggested text could be removed if any reason to doubt GW151226 is found.

The fact that this event was stronger than typically expected for a first detection and of a type (a binary black hole system) that had large uncertainty in the expected rates opened up a number of investigations that might not have been possible with a binary neutron star coalescence at lower signal-to-noise ratio. The collaboration has responded well to this opportunity by producing a number of results spanning astrophysics, gravitational wave predictions, and tests of general relativity.

That said, significant material is still being consolidated in the main detection paper. As of the writing of this report, there are some aspects that are under review or that have completed their review too late for this committee to independently evaluate their completeness. We have no reason to doubt the correctness of any of this material, and none of these elements represents any danger to the core detection claim itself. However, because this discussion will be new in v9, we recommend that the vote on the paper be delayed to give the collaborations the chance to discuss the proposed final draft fully. A delay of 2-4 days to give more time for everyone to become comfortable with all of the material in the paper should not unduly impact the overall timeline.

We recognize that the companion papers which are supposed to support the main detection paper are still in varying states of maturity, with some recent results being checked for accuracy. We recommend that the collaborations make and announce a clear policy about how to deal with possible changes to numbers that might arise during these final checks. It is important that we minimize the discrepancies between numbers in the detection paper and the companion papers. We recommend that we use the close of the vote on the main detection paper as a deadline for discretionary changes to the material in the companion papers. Discretionary changes include new results and new enhancements of existing results, as distinct from changes to correct errors, which we should always do.

Detailed Questions:

- Have the i) analyses pipelines and associated software … been checked and certified to produce trustworthy results?

Some of the reviews of the analysis software and searches have been delayed and are only recently available for our evaluation. Furthermore, some of the reviews have been completed without full agreement by the search groups on the results. We have no reason to doubt the validity of anything which has been cleared by the reviews, but we have not been able to independently evaluate them.

In this regard, the fact that multiple searches have found GW150914 with high significance is greatly reassuring. The consistency of the results from different pipelines is remarkably good, and by itself is a strong indication that there are no major errors in any of the searches. We believe that the agreement among independent searches is sufficient "proof" that the results can be trusted.

That said, there are some disagreements between searches in the detailed estimates of false alarm rates, etc. For example pyCBC and gstlal have different methods of estimating false alarm probabilities which have not been completely reconciled. In such cases, we recommend that the detection paper rely on the more conservative numbers and on the pipelines for which we have the greatest experience. Newer pipelines should continue to be pursued because they may have significant advantages once they are more fully vetted.

The Detection Committee has given some feedback to the Paper Coordinating Team concerning the clarity of the way the search results are presented in v8 of the paper. Like the rest of the collaboration, we are eagerly waiting to see v9 of the paper to see the response to these comments.

- Have … ii) the detector control software been checked and certified to produce trustworthy results?

It should be noted that an "ordinary" malfunction of control software, such as a lost data sample or A/D overflow would not lead to the distinctive waveform recovered for GW150914 in either interferometer. In addition, unless such malfunctions are correlated, e.g., triggered by a condition explicitly dependent on GPS time, their coincident occurrence to a degree detectable in the H1 and L1 DARM signals can be accurately estimated from time slide studies and hence are automatically included in existing false alarm probabilities.

One should worry more about GPS-triggered events in the real-time system, inserted either inadvertently or maliciously. Concerning the former, one should also remember that following a fiber breakage in the H2 interferometer in March 2012, provoked by a D/A value going out of a safe range, an extensive review of the CDS real-time software system was carried out [LIGO M1200346]. A number of recommendations were made and then implemented to improve overall reliability, security and traceability of CDS software. These safeguards have included more formal vetting of real-time software changes, more restricted access to CDS computers, and adoption of svn for version control.

The real concern then is an accidental or malicious injection of a simulated astrophysical signal. One natural mode of injection would be via the hardware injection infrastructure [T1400349] in which signals are injected into excitation channels, which are added to the input D/A channels of the electrostatic drives of end mirrors or of the photon calibrators for the end mirrors. Accidental and intentional injections using this infrastructure have been ruled out, however, by careful and detailed examination of all excitation channels (unblinded and blinded), [EVNT entries 11195, 11253, 11331], as discussed below under malicious tampering.

In addition, the CDS team has examined in a comprehensive audit the conditions of the DAQ and front-end computers at the time of GW150914 [LIGO T1500541 and EVNT entries 11201, 11260, 11266, 11303, 11309, 11318, 11320, 11329, 11330, 11430 and 11431]. No evidence of suspicious logins or processes was found on CDS gateways or on injection computers.

One potentially serious worry specific to detector control software, however, is a malicious injection carried out directly into the A/D channels for the OMC photodiode pairs used to read out the interferometer GW signals. In principle, a properly scaled waveform added to those A/D channel contents at times within 10 ms of each other for H1 and L1, could mimic a true GW signal, without abnormally affecting other servo channels. Such an addition of a waveform would require either that 1) the front-end model for the OMC photodiode readout be modified to perform the injection; or 2) a separate, rogue process be run in parallel and modify A/D contents on the fly without disturbing the servo perceptibly.

The CDS team's audit considered possibility 1) explicitly. The models running on all of the front-end machines were compared to compilations of archived programs, modification dates checked, and all unexpected discrepancies tracked down and resolved. For the critical OMC front-end code, line-by-line code examination was carried out, with no evidence of a waveform injection or other tampering found. The most recent changes to the OMC code prior to GW150914 were examined, in particular, and were consistent with the purpose of those changes (calibration model updates).

Possibility 2) -- a parallel rogue process -- is more difficult to rule out, based on forensic evidence alone, because it would presumably be designed to remove any signs that it had ever been run. The CDS team notes in T1500541 that "logging methods were not in place at the time to verify that additional, unauthorized code was not running on the RT [real-time] computers at the time that might affect the data." Such a RT process, however, would face a substantial technical challenge in intercepting the A/D channel contents and adding signal waveform contents to them at 16,384 Hz while not interfering with the front-end model code and causing perceptible symptoms of time jitter, which are not seen in GW150914. However, the logging methods noted by the CDS team above were in place for the second event GW151226, and all object files loaded at that time were under configuration control, the machines having all been restarted between Sep 14 and Dec 26 in association with a site-wide power outage (LHO) or Tuesday reboots (LLO), adding confidence that no parallel rogue process exists.

In summary, there is absolutely no evidence to support any concern that detector control software was used, inadvertently or deliberately, to inject the GW150914 signal. While that possibility cannot be ruled out with complete certainty, because of the possibility of a self-destructing rogue front-end process, this committee finds that scenario utterly implausible, not only because of the daunting technical challenge, but also because it would require a nefarious conspiracy among highly trusted colleagues.

> ***- Is the data around the time of the detection of sufficient quality and the instrument behavior sufficiently understood that we can confidently rule out an instrumental artifact?***

Yes.

The interferometer performance was clean and well understood, by every measure available:

- Both H1 and L1 were in the middle of multi-hour locks. Calibration was steady; PSD was steady. All measures shown on Summary Pages' FOM plots were steady. Instruments timing systems were in sync.

- Signals match in both OMC photodiodes at each interferometer. The signal was really on the light, as expected. A trace through the locking loops at each interferometer agrees.

Checks for possible instrumental or environmental disturbances showed that there were none that could have caused the event:

- No system excitations, including hardware injections, were active at the time of the event.

- No veto segment overlapped the event. No auxiliary channel showed any glitches strong enough to suggest an instrumental artifact.

- Environmental sensors were carefully checked. There's a large margin between the environmental disturbances and the level that would have been required to cause a signal as strong as GW150914. A cosmic ray detector at H1 showed no coincident event.

We made special checks for all known classes of instrumental transients. None were present at the time of the event.

- The characteristic signature of a scattered light disturbance is not present at the time of the event. Nor does the event look like any of the other classes of glitches known to affect the interferometers.

- DAC glitches weren't present in the instrument at this time. There is no sign of any saturations or overflows that could have made a spurious signal.

Similar checks were also made for G197392, although it was neither possible nor appropriate to perform those checks as thoroughly as for GW150914. No aspect of instrument performance disqualifies G197392 as a potential gravitational wave signal. We do note, however, that the data were not as clean at that time as for GW150914; glitch rates were higher overall, and there is a low frequency glitch in the L1 data at the same time as the event. We feel that these facts support an attitude of caution in interpreting G197392.

- Is the case strong enough to claim a detection? Is the statistical confidence sufficiently compelling in your estimation?

We are convinced the case to claim GW150914 as a detection is very strong. An environmental or instrumental artifact has been ruled out (see previous section) and the statistical significance is definitely compelling enough for us to be confident that GW150914 is a gravitational wave signal from a BBH source.

The signal has been consistently identified as an H1L1 coincidence and ranked as very significant by several search pipelines, both those that target CBC signals and those designed for generic transient signals. (We have fulfilled the LVC's multiple pipeline policy many times over!) . Bayesian parameter estimation tools also overwhelmingly favor the coherent signal hypothesis over the (Gaussian) noise hypothesis. Consistency tests between the observations and GR expectations haven't revealed anything unexpected. All the information forms a consistent picture of both the astrophysical nature of the signal and the agreement of the observed signal with that expected from a BBH merger characterized by the inferred parameters.

Specifically: pycbc estimates a FAR of less than 1 in 200000 years (>5.1 sigma). gstlal estimates a FAP of 1.4 x 10^{-11} (6.6 sigma). cWB estimates a FAR of < 1/22500 years (> 4.6 sigma) or 1/8400 years (4.4 sigma). The CBC Bayesian analysis provides a very large Bayes factor ln B_s/n ~290. The false alarm probability reported by each search comprehensively takes into account the look-elsewhere-effects involved in that search, so that we believe it is sound to base a detection claim on a significance in excess of 5 sigma.

We have given feedback to the Paper Coordinating Team aimed at improving the clarity of the presentation of the search confidence level. However, we have not seen v9 and cannot comment on this aspect of the detection paper.

- Is there additional evidence to support the detection claim, eg, from other astrophysical observations? If not, is the lack of additional evidence detrimental to the detection claim?

We are convinced that a secure detection claim can be made on the basis of the gravitational wave data alone. The claim that the transient is due to a BBH merger is supported by the parameter estimation results, both those that use CBC templates and those that don't (from the Burst analysis). The BBH claim is further strengthened by the fact that, with minimal assumptions, the Burst Group's reconstructed waveform is in excellent agreement with the CBC-based waveforms.

Sky maps were circulated to EM partners (about 3 days after the event), but to our knowledge, no interesting EM counterpart, even remotely associated with GW150914, was ever identified. However, this negative EM result is no cause of worry about our detection claim. The reason is that the Event is robustly identified as a BBH (no support is found in parameter estimation towards masses that might be close to NS). For BBH mergers, no reliable/robust predictions for EM counterparts exist, as no matter is expected to be involved in such mergers, and interstellar medium densities are too low by many orders of magnitude for any significant EM production. Even somewhat exotic models in the literature that discuss potential EM counterparts for BBH mergers, when scaled to the GW150914 masses, lead to effectively undetectable signatures in any EM band.

Given the robust BBH nature of GW150914, the BH masses involved, its distance at few-several 100's of Mpc, the latency in our communication to the EM partners, the size of the GW source position error boxes and the lack of full coverage from EM observations, and the lack of predictions for any detectable EM emission, we are confident that the absence of any observed EM counterpart is in no way detrimental to our detection claim.

- Are we confident that the event is not the result of malicious tampering?

We have gone to great lengths to investigate the remote possibility that GW150914 (the Event) may have been placed in the LIGO data through malicious tampering, either as a hardware or software injection. This has been studied in detail, and our conclusion is that we can rule out any serious possibility that the Event was a result of malicious tampering. While not impossible, it would require an internal conspiracy involving a group of experts with very disparate expertise and access to deep parts of the instrument and data. This has been deemed highly implausible.

Malicious tampering would occur in the form of injections, where a fake signal is introduced into the data.

We note that there are times when legitimate “normal” injections are made. During previous science data taking runs fake signals have been injected into the LIGO data (either via hardware — where a force with an astrophysical waveform is applied to the interferometer mirrors, or via software — where an astrophysical waveform is introduced into the data stream). At the time of the Event, no such deliberate “blind” injections program was active. Even so, an extensive study of hardware excitations channels showed that no deliberate or accidental injections were made.

This brings us to the question of rogue injections. Three scenarios that cover a range of ways in which malicious injections might be made have been considered:

(i) Frame spoofing. This involves replacing the original data frames with fake ones that contain an event that was not in the real data.

(ii) Double blind injections. These injections are introduced in ways that are not recordable. They may be injected through an excitation channel whose state is not recorded, or via a path that is not normally used for injections such as the photon calibrator, or as an Easter egg hack through modification of the digital system. (Easter egg hacks involve an undocumented feature in the software that causes an unexpected chain of events to occur.)

(iii) Analog hacks. This would involve adding hardware that changes the signals in analog before they are digitized, so not recording by the data and/or controls computers.

Frame spoofing. The gravitational wave signal is imprinted on the light, converted into an analog signal on the photodetectors, digitized and sent to the CDS (Control and Data Systems) computers. The front computers also pass this signal to LDAS (LIGO Data Analysis System) for storage and distribution. The LDAS data is quickly distributed from the observatories to remote storage and analysis sites, so any software malicious injection would require working at the observatories or access to multiple complicated computing systems. For this reason frame spoofing would have to be done simultaneously on multiple servers at LHO, LLO, and Caltech, while erasing all traces of access and activity, such as remote log files. This would not be impossible to do, but what makes it extremely implausible is the complexity of the signal chain. More than 30 versions of the DARM signal are stored at various points along the signal chain. To reproduce these signals, and the filtering between each version, a hacker would need to understand both the instrument signal chain, and the CDS servers, configurations files and live parameters deeply. We do not believe that these capabilities reside in any single known individual, or even a small group that might collude. Frame spoofing has, therefore, been ruled out.

Double blind injections. Again the redundancies of the signal chain help rule out the possibility of malicious injections being added. If they are added along the normal DARM chain, they will be recorded in multiple channels, so easily found. For example, one expects to observe excess counts in the electrostatic drive channel, as is confirmed for known hardware injections using ESD actuation. If they are injected through another path, e.g. PCAL, the injection should leave a large transient in that channel that would be coincident with the DARM trigger, including the readout channel for the photodiode that witnesses the reflected PCAL beam. To suppress the appearance of that transient in the auxiliary channel data would require inverting the injection, which in turn would require knowing the exact coupling to DARM. No such inverse filters were in place at the time of the Event. Furthermore, all auxiliary channels, including all excitation channels, have been checked exhaustively for coincident transients, with nothing significant at the time of the Event. What was found is that all channels that should contain the Event signal do, and all channels that should not, do not. [LIGO G1501282, T1500536 and EVNT entries 11273, 11308, 11336]

Similarly, an Easter egg hack would require the same kind of generation of fake signal and, in most insertion points, actuator inversion to hide it, but also CDS administrator privileges at both observatories, deep knowledge of the CDS system, and the ability to remove the malicious code before the post-Event inspection. The special case for which mechanical actuator inversion is unnecessary is discussed under detector control software above. All checks of the code versions have not shown any unusual activity or code changes at the time of the Event, as discussed there.

Analog hacks. There are a few analog electronics modules where a signal could be injected into a mirror actuator after the monitor points that are recorded. This would require embedding a new piece of hardware (perhaps as rudimentary as an iPod nano and some clip leads) into the actuator driver modules. To pull this off, a hacker would need access to the physical modules, in-depth knowledge of the electronics and signal chains, programmable injection devices with millisecond resolution at both sites, and of course the right astrophysical waveform to inject. The disparate types of expertise needed alone might rule out this possibility. But just to be sure, a thorough inspection of the electronics modules was carried out some days after the Event, with nothing untoward noticed. Card reader access and environmental monitors showed no unexpected activity. Furthermore, individual chasses were opened up at LHO to look for rogue devices. None were found. A similar process will be carried out at LLO shortly after the completion of O1.

The above is a summary of numerous studies that are described in dozens of EVNT logs, and summarized in LIGO-L1500138 and LIGO-G1501317. Access at the sites was checked via card reader logs, work permit logs, physical inspection, physical environment monitors, and reports of personnel on-site.

> ***-Is the presentation of partial O1 results compelling without presenting the rest of the O1 analysis in progress?***

Yes. GW150914 is the loudest event in all of O1 data from chunk 1 to chunk 8 (which goes up to 27 December). Using 16 days of coincident data, the pyCBC pipeline assigned a FAR limit of 1 in 200,000 years, or a significance that is far higher than 5.1 sigma. In fact, under the assumption of a noise model that the data contains no gravitational wave events, i.e. that GW150914 is a background event, the gstlal method gives a point estimate of the false alarm probability of this event to be 1.4e-11 or a significance of 6.6 sigma (these numbers are being finalized at the time of writing but their exact value will not change any of the conclusions below). Increasing the time analyzed could improve the estimate of the significance, but this benefit is marginal, in our view.

The second loudest event in 16 days of data considered for the publication of the discovery paper has a FAR of 1 in 2.3 years and a FAP of 0.02 (~2 sigma). This is not a significant enough event to claim a detection on its own but both time-frequency maps and parameter estimation runs are consistent with what is nominally expected for a real event and it definitely adds value to the detection paper.

The greatest advantage of including all of O1 data so far is not the decreased FAR of GW150914 but the Boxing Day event. As reported by the CBC analysis, both pyCBC and gstlal detect GW151226 with the same template, assign (inclusive) FAR of 1/9600 years (3.3e-12 Hz) and 1.2e-12 Hz and FAP of 1.4e-6 and 6.4e-7. So what brings an even greater confidence to a sceptic is GW151226. As discussed in the summary we recommend that GW151226 be mentioned explicitly, with appropriate cautions about its preliminary status.

-Is the physical (GR) and astrophysical interpretation of the data appropriate?

As noted above, the nature of this event, a binary black hole coalescence with relatively high signal to noise, opened a number of physical and astrophysical studies that would not have been possible in what many people might have imagined for the first detection. The collaborations have responded well to this opportunity to perform a large number of studies that are being documented in a number of companion papers.

The Detection Committee has reviewed these companion papers with a limited scope: are the results that are shown in the companion papers consistent with those shown in the main detection paper? and are they in any way (tone, substance, implications) in conflict with the main messages of the main detection paper? Our official communications to the authors of the companion papers have centered on these questions only. We have not as a committee reviewed the detailed technical content of the companion papers for correctness, though many of the members of the Detection Committee have sent detailed comments to the authors of the companion papers either as individuals or in other roles (e.g., paper reviewers).

A subset of the companion paper results has been selected for mention in the main detection paper, though obviously the discussion of these points is greatly simplified in the main detection paper. The discussion in the section “Source Discussion” contains the majority of these items. A large fraction of our time has been spent in discussions about the appropriateness, the accuracy and the clarity of the physics and astrophysics results which are in the detection paper.

In general, we agree that the items which have been selected are appropriate. However, we had some comments concerning the presentation of several of the items, which we have given the Paper Coordinating Team since the release of v8. We understand that many of these suggestions will be addressed in v9, but we cannot comment on them at this time.

Appendix

Charge to the Committee

To: Stan Whitcomb, Frederique Marion (co-chairs) and members of LSC-Virgo Detection Committee

From: Gabriela González and Fulvio Ricci, LSC and Virgo spokespersons

The Spokespersons and other appropriate parties have agreed that the detection case for GW150914 is in the appropriate stage to move to Step 3 of the Detection Procedure LIGO-M080010-v2, “Preparing for the decision”.

As part of this step, we formally charge the Detection Committee to objectively examine the detection case as presented in LIGO-P150914. We have asked the Paper Coordinating Team (chaired by Peter Fritschel and Eric Chassande-Motin) to help you liaise with the appropriate experts of each section.

At the end of your investigations, we’d like a brief report summarizing your findings. The report will be made available to the LSC and Virgo Collaboration when making the decision to submit the paper for publication. Drafts of companion articles that are intended to be public at the time we announce the detection (after having the paper accepted for publication) are in https://wiki.ligo.org/DAC/GW150914papers.

We’d like either a report or an estimate of the delivery time the report will take by January 18. We know this is a tight timeline, but you shouldn’t compromise your investigations in any way, and take the time you need.

We’d like you to challenge the detection claim asking questions that skeptical scientists may ask of the LVC when reading the detection paper and companion documents. The following is by no means an exhaustive list, but questions you should address in you report include:

- Have the i) analyses pipelines and associated software and ii) the detector control software been checked and certified to produce trustworthy results?

- Is the data around the time of the detection of sufficient quality and the instrument behavior sufficiently understood that we can confidently rule out an instrumental artifact?

- Is the case strong enough to claim a detection? Is the statistical confidence sufficiently compelling in your estimation?

- Is there additional evidence to support the detection claim, eg, from other astrophysical observations? If not, is the lack of additional evidence detrimental to the detection claim?

- Are we confident that the event is not the result of malicious tampering?

- Is the presentation of partial O1 results compelling without presenting the rest of the O1 analysis in progress?

- Is the physical (GR) and astrophysical interpretation of the data appropriate?

You should view your role as somewhat adversarial, challenging the detection claim. You should feel free to interact with the paper Coordinating team, the teams assigned by them to defend statements in the paper, and the analysis and review groups to request further checks and actions as needed.

We stress that the role of the Detection Committee is purely evaluative. The ultimate decision about whether or not to make a detection claim and submit the paper for publication will be made by the LSC and Virgo Collaboration. Therefore the report should address the validity of the claim along the lines of “We believe (or don’t believe) that the analysis performed by groups in support of the detection claim is valid (not valid) and substantially correct (or incomplete or flawed).”

Thanks in advance for your work in this very important process for our field,

Gaby and Fulvio.