A-level

Philosophy

7172/1 Epistemology and moral philosophy

Report on the Examination

7172
June 2024

Version: 1.0

Further copies of this Report are available from aqa.org.uk

Introduction

What follows is a question-by-question commentary on the key trends in the performance of students on component 7172/1: Epistemology and Moral Philosophy. In the course of this commentary, reference is made to student responses, the question paper, the assessment objectives, AQA's *A-Level Philosophy (7172) Specification* , and the Mark Scheme. In compiling this report, the observations of the lead examiner have been supplemented by the evidence provided by senior examiners and their team members. Reference is also made to statistical evidence to provide contextualisation to the qualitative judgments and findings outlined in this report. We are now advised against using precise statistics (eg mean averages for each item) given the inevitable (albeit small) variations between awarding and results day. We will therefore refer to statistical trends in more general terms, offering provisional and approximate data on the distribution of marks for each item.

Summary Findings

The (mean) average performance this year is fractionally over 53%, down by a mark on 2023. That is largely explained by unusually high scoring 12-mark questions last year. The outcomes on Epistemology and Moral Philosophy are very evenly balanced this year, with students doing fractionally better on Epistemology. On average, students were able to access over 50% of the available marks on both sections of the paper. This supports the view of the senior marking team that this paper was accessible to the majority of students. Many students demonstrated a sound knowledge of the key concepts, arguments, and theories, and were able to use this to score well on AO1. On the AO2 (essay) questions, the vast majority were able to argue with some degree of intent and to present their answers in the form of an argument. Whilst inevitably there were some students who did not have the knowledge or skills to address certain questions effectively, the majority were well prepared and there were no surprise questions that tripped them up or sent them in the wrong direction. As such, there were very few instances of items going unattempted by students (less than 1% overall), and with fewer students scoring zero marks for items compared with last year. The majority also had a good sense of how to answer the different question types in terms of timing, skills, and focus.

As noted already, the paper was very well balanced in terms of student outcomes. Some noteworthy differences came on the 5-mark questions, where students tended to perform better on Epistemology, although this was evened out by slightly better performances on the 3 and 12-mark Moral Philosophy questions. It was encouraging to see students performing equally well on both 25-mark questions and to see an improved performance on these essays compared to previous years. The mean averages for Question 05 and Question 10 are just below 12 marks, with Epistemology scoring narrowly higher.

General Observations

Senior examiners made a number of general observations this year, which will be elaborated on in the main body of the report, but which can be usefully summarised as follows:

1. Where questions are testing a student's knowledge and understanding (AO1) of the particular argument, distinction or principle of a named philosopher, students are reminded that they need to follow, as closely as possible, the logic and detail of the argument as presented by that philosopher in the original text. These are referred to in the 'Set Texts' sections of the *Specification* and are the primary sources. This can be found on page 9 of the *Specification* under

‘Subject content’ where it states: ‘Students must demonstrate an understanding of, and the ability to make a reasoned evaluation of, the arguments set out in those texts’. See questions 03, 07 and 08 for more discussion of this.

2. Whilst familiarity with logical terms and methods is improving, it could be improved further – eg abductive reasoning (see questions 03 and 08), drawing the right/appropriate conclusions, and using Mark Scheme language with sufficient justification – eg ‘stronger/weaker’ and ‘crucial/less crucial’ (see questions 05 and 10) – in the content of the essays.
3. On 25-mark questions: (a) Some students tried to cover all the material from the *Specification* (and sometimes went beyond it) in their answers, which often resulted in a lack of detailed analysis and robustness of evaluation. Students would often be better served by being more selective, covering less material but in greater depth. (b) Some students appeared to have learned essay plans by rote but without clearly understanding all of the content or the logic. This meant that some formulaic responses lacked coherency, especially when the student attempted to engage in the evaluation of material that they had not fully grasped (see questions 05 and 10 for more on this).

Assessment Objectives

AO1: Demonstrate knowledge and understanding of the core concepts and methods of philosophy, including through the use of philosophical analysis.

AO2: Analyse and evaluate philosophical argument to form reasoned judgements.

Section A: Epistemology

Question 1: What is meant by (a) a necessary condition and (b) a sufficient condition? (3 marks)

General Observations

This question tested students’ knowledge and understanding (AO1) of two fundamental concepts regularly employed in philosophy when attempting to provide a precise definition of a concept or phenomena. They are referred to in the *Specification* - ‘What is knowledge?’ (3.1.1) – under ‘The tripartite view’. Whilst the concepts are explicitly referred to in the Epistemology section, given that they have such broad philosophical scope, students did not need to explain them exclusively in epistemological terms, although many students chose to give epistemological examples, particularly in relation to the tripartite definition of knowledge or some variant. However, students were able to explain the conditions in numerous ways (eg metaphysically or logically) and any legitimate way of explaining how the conditions operate (often through the use of relevant examples) was credited as long as students demonstrated that they understood the key points – ie that ‘a necessary condition of X is one that is *needed/required* in order for X to be the case, whereas a sufficient condition of X is one that is *enough*, if satisfied’.

Students were slightly less effective in meeting the demands of this question than on the corresponding 3-mark question on the Moral Philosophy section. However, the most commonly awarded score was still 3 marks (almost 39% of students); almost 23% of students scored 2 marks; and over 28% scored 1 mark. This meant that almost 90% of students were able to write something worthy of credit.

Where students performed well

To score full marks on this question, students were required to provide a clear and precise definition of each term. Many did so in terms of a condition being '*needed* (necessary)' or a condition being '*enough* (sufficient)'. There are a number of examples in the indicative content of the Mark Scheme that capture this and other approaches. Some answers simply gave the definitions with no examples, whilst others chose to illustrate their answer, often of both conditions but sometimes only of one. Each approach was acceptable - as long as the answer was correct and the examples relevant - and these students were awarded full marks accordingly. When used effectively, examples do not constitute redundancy and they are often helpful for students when explaining/showing an understanding of key concepts. However, there are some potential pitfalls that are mentioned below.

Common / noteworthy weaknesses in student responses

There were a number of reasons why students failed to gain full marks. These ranged from only showing a clear understanding of one condition (typically the necessary condition) and a partial or imprecise understanding of the other, or there might be a lack of detail and/or precision on both conditions although the answer was still substantively correct (typical of 2-mark responses). Alternatively, some showed imprecision in the answer because of an intrusive example which compromised the accuracy of the answer overall. There are numerous examples listed in the indicative content of the Mark Scheme that illustrate these types of response.

The most common errors of understanding included straightforwardly confusing 'necessary and sufficient conditions' with 'necessary and contingent truths', or incorrectly claiming that 'both necessary and sufficient conditions were required/needed' when only necessary conditions are. For example, getting 50% in an exam might be *sufficient* for passing the exam but it is *not necessary* as you can pass the exam by getting a different percentage (eg by getting 51%) and therefore getting 50% is sufficient but is not required/needed. It may, however, be said to be necessary to get *at least* 50% to pass.

Question 2: How does the argument from perceptual variation present an issue for direct realism? (5 marks)

General Observations

This question required knowledge and understanding (AO1) of how a key philosophical argument presents a challenge to direct realism, one of three named theories from the 'Perception as a source of knowledge' section of the *Specification* (3.1.2). This was the best performing question on the paper with the mean average significantly higher than that of the paper overall. The vast majority of students (over 85%) managed to score at least 3 marks on this question, showing that they had grasped the substantive content of the issue. In order to do this, they needed to demonstrate an understanding of both direct realism and the argument from perceptual variation, and then some understanding of how the argument impacts the central claims of direct realism. Typically, students gave relevant examples to illustrate the argument, with the most common being 'Russell's table', followed by Locke's example of 'lukewarm water' or some of Berkeley's examples. Over a quarter of all students scored the maximum 5 marks. At the lower end, over 99% of students were able to score at least a mark on this question, with over 11% awarded 2 marks.

Where students performed well

At the upper end of the performance scale, students who grasped the substantive content were then able to go on to give fuller, clearer, and more precise accounts, which demonstrated a deeper understanding of the argument and how exactly it impacts the claims of direct realism. Some set out the main claims of direct realism – 'that mind independent objects exist and that we perceive them directly with no intermediary'. They then set out the argument from perceptual variation (sometimes in its logical form but not always) before identifying the main challenge – ie that, *pace* direct realism, 'we cannot be directly aware of mind-independent objects because, whilst our perception of the object and/or its properties changes depending on our perspectival position, the object itself does not change'. Some then went on to draw the conclusion that this meant that what we are directly aware of is a subjective intermediary like sense-data, but students did not need to do this to gain full marks.

Common / noteworthy weaknesses in student responses

Responses that gained 2 marks lacked the substantive content, giving some relevant points but not precisely, so not enough to demonstrate sufficient knowledge and understanding of how the argument presents an issue for direct realism. For example, some students were able to give a generally correct (although imprecise) account of direct realism - eg a crude definition of naïve realism - and identify a relevant example but without linking the two. Others gave a generally correct but imprecise account of the argument but failed to discuss direct realism. Another common error was to use examples from other arguments to illustrate the issue, the most common being the 'crooked pencil' from the argument from illusion, or to mix up the arguments – eg by explaining the argument from illusion or hallucination rather than the argument from perceptual variation.

Question 3: Explain how Bertrand Russell responds to scepticism by arguing that the external world is the 'best hypothesis'. (5 marks)

General Observations

This question required knowledge and understanding (AO1) of Russell's response to external world scepticism from the 'Perception as a source of knowledge' section of the *Specification* (3.1.2). Students performed well, with the mean average well above that of the paper as a whole, and it was the second best performing 5-mark question. Over two thirds of students managed to score at least 3 marks on this question, showing that they had grasped the substantive content of Russell's response. In order to do so, they needed to demonstrate an understanding of the sceptical challenge and Russell's response to it, with better responses showing an understanding of the non-deductive nature of Russell's argument and some grasp of the competing hypotheses. Typically, students went on to use the example of 'Russell's cat' and better answers used this effectively to explain why the external world is the best hypothesis. Those that did not use an example often struggled to give a clear explanation. Indeed, it is hard to explain Russell's claim in the abstract, and so students really benefited from using examples – as they so often do. Almost 15% were awarded the full 5 marks. At the lower end, over 95% of students were able to score at least something on the item, with just over 17% of students awarded 2 marks.

Where students performed well

At the upper end of the performance scale, students who grasped the substantive content were able to go on to give a fuller, clearer, and more precise account, which demonstrated a deeper understanding of Russell's response. In particular, these students identified the abductive nature of Russell's argument. Some examiners found it surprising that more students did not do so, given that this is one of the key terms that students are supposed to be able to demonstrate an understanding of (see pages 9-10 of the *Specification*). An ability to employ key philosophical vocabulary, including the use of logical terms, is one way that students can add precision to their answers. Another feature of higher-level responses was when students mentioned 'that belief in the external world leads to no inconsistencies,' and is an 'instinctive' belief, with some accurately explaining the significance of this claim for Russell's argument and the conditions under which we should accept an instinctive belief (i.e. in the absence of positive reasons to doubt it). However, this was quite rare, and students could still gain full marks without mentioning these points.

Common / noteworthy weaknesses in student responses

Whilst most answers identified Russell's claim that the external world hypothesis is 'simpler', weaker responses were unable to explain this well, claiming that the external world is the best hypothesis because it is 'easier to believe' or because it 'makes sense' without unpacking what might be meant by these expressions. Another common error was using examples and arguments that were more reminiscent of Locke than Russell – eg concerning the 'involuntary nature of perception'. Whilst imprecise, this was not penalised if students were able to use additional material to help them develop the Russellian argument. In some cases, however, there was either too much blurring and/or not enough specific detail on Russell. Such answers often failed to provide the substantive content, giving some relevant points but not precisely, and so not demonstrating a sufficient understanding of Russell's response.

Perhaps the most common error was when students confused explanatory simplicity with ontological simplicity and invoked 'Ockham's razor'. The misuse of Ockham's razor is a common error seen both in this and previous series. Although it is of course relevant epistemologically, since it arises in the course of trying to understand or explain something, Ockham's razor is a principle of ontological parsimony, which aims to prohibit the unnecessary multiplication of entities in our explanations. If phenomena can be explained using fewer entities or types of entities, then that explanatory hypothesis should be favoured over competing hypotheses, assuming that the preferred hypothesis explains the phenomena adequately. For example, idealists could claim that their theory is ontologically simpler than indirect realism given that it only posits the existence of one ontological category (ie the mental) whereas indirect realists posit the existence of two (both physical and mental). Therefore, if it can be shown that idealism adequately explains perceptual phenomena, then, by Ockham's razor, it should be favoured over indirect realism because it is both explanatorily adequate *and* ontologically simpler. As an indirect realist, Russell's argument invokes physical objects, perceivers, and sense-data, which is less ontologically parsimonious than the alternative hypothesis that Russell presents (which invokes perceivers and sense-data but no mind-independent physical objects). As such, it is incorrect to claim that Russell's response is one that conforms to Ockham's razor and Russell does not use this principle as part of his argument in the set text *The Problems of Philosophy*.

Question 4: Explain Descartes' cogito <u>and</u> an empiricist response to it. (12 marks)

General Observations
This question required knowledge and understanding (AO1) of one of the most famous claims/arguments in modern philosophy - Descartes' cogito – and (less famously) an empiricist response to it from the 'Reason as a source of knowledge' section of the *Specification* (3.1.3). It tested students' ability to not only explain both parts of the question, but also to logically integrate these into a coherent response. Overall, this did not perform as well as the corresponding 12-mark Moral Philosophy question, but both averaged comfortably over half marks (6), and a majority scored 7 or more on this question. Nearly a quarter of students were placed in the top level (10-12 marks), while just under a third were placed in the 4-6 band. At the lower end, almost 19% of students scored in the 1-3 band. Under 2.5% of students failed to pick up any marks.

To access 7 marks or above, it was essential for students to provide at least some relevant and correct material pertaining to both parts of the question and so to demonstrate the skill of integration. A response that only addressed the requirement to explain Descartes' cogito but did not provide any relevant material on an empiricist response (or vice versa) could not demonstrate the level of detailed understanding and integration that is required for marks higher up (ie 7 and above), so could at best be described as 'largely correct' with 'most points made clearly'. In the vast majority of cases where students offered only a partial response, they were able to say something about Descartes' cogito but tended to struggle with an empiricist response. A student could still gain up to half marks for a detailed and precise explanation that showed evidence of logical linking on just one part of the question.

A significant number of students opted to discuss the cogito in the wider context of the three waves of doubt and these responses were variable in quality. Other students went straight to the task of explaining the cogito. Either approach was acceptable although better responses tended to get straight down to business and so saved valuable time. Of those who opted to explain the cogito in the wider context of the 'three waves of doubt', some did so pithily and to good effect, while others gave more extended treatments which still got them to where they needed to be to address the most salient issues. Some students, however, seemed to be using the three waves of doubt as a cover for not being able to explain the cogito in sufficient detail.

In some cases, students presented the cogito as an 'a priori intuition' and were able to unpack these terms precisely, explaining why Descartes presents it as such in the *Meditations*. Others opted to present the cogito as an argument, and this allowed them to explain in more detail certain aspects of Descartes' reasoning. Both approaches were welcome, and the Mark Scheme acknowledges this. How students chose to present the cogito – either as an intuition or an argument – impacted how they presented the particular empiricist response they chose to address, and so students needed to ensure that they were able to integrate them correctly and coherently.

Where students performed well

Students who performed well generally approached this question in two parts and then linked them together at the end. This worked effectively, with the best answers not only explaining both parts of the question but also showing exactly how their preferred empiricist response impacted Descartes' cogito. These students were able to explain Descartes' cogito – eg by linking it to the notion of 'intuition' and 'clear and distinct ideas' and then giving detailed and precise explanations of these terms and how they relate to the cogito. As stated above, some did this in the context of the three waves of doubt whereas others chose not to. Having explained the cogito, they then went on to explain an empiricist response, and these tended to be one of Hume's objections, either his 'bundle theory' objection and/or by utilising 'Hume's Fork' (some of the best responses were able to integrate both of these Human responses - see the Mark Scheme for details on how this could be done), or Russell's 'persisting self' objection.

Common / noteworthy weaknesses in student responses

Students who performed less well did so for a number of reasons, some already outlined - eg by only addressing one part of the question and/or by spending too much time contextualising the cogito and not enough time actually explaining it. In a similar vein, some students wanted to write down everything they knew about Descartes from the three waves of doubt to his proofs for God's existence and his proof of the external world. Again, this often masked a lack of knowledge and understanding of the particular demands of the question, constituting redundancy and a lack of sufficient relevant material.

It is notable that a significant number of students struggled with the empiricist response, writing either nothing on this or giving very brief, tangential, incoherent, or incorrect responses. Middle range responses (7-9) that did demonstrate knowledge and understanding of an empiricist objection, which tended to be based on Hume's Fork, were able to address both parts of the question and provide some link between them, but they often struggled to integrate them effectively – eg by arguing that Descartes' cogito is 'a synthetic a priori claim but, given Hume's Fork, there could be no synthetic a priori truths, and so Descartes' cogito is false' - but without explaining this in detail or precisely so as to make the impact of the empiricist response on the cogito fully clear.

Question 5: Is there a successful way in which propositional knowledge can be defined? (25 marks)

General Observations

This was the first of two questions on the paper designed to test both A01 and AO2, with the majority of marks for A02. The question centred on evaluating competing definitions of propositional knowledge, which is the focus of the topic: 'What is knowledge?' (*Specification* 3.1.1). Statistically, this was the second hardest question on the paper, with a mean average well below that of the paper as a whole, but with students averaging just under 12 marks. This was higher than on the corresponding Epistemology question in 2023 (10.7 on philosophical scepticism), and there has been a significant improvement overall since we returned to examinations after the pandemic. Almost a quarter of students accessed the top two bands this year, scoring at least 16 marks (this compares with just 17.14% at the same time in the series last year). At the very top (21-15) there was also upwards movement, with over 5 % of students achieving this. The most frequently accessed band was 11-15 (at almost 40%). At the lower end, just over a quarter of students scored in the 6-10 band. The bottom scoring band (1-5 marks) was awarded to just over 8% of students. Less than 3% of students failed to pick up any marks.

In general, responses were quite formulaic, although the topic can lend itself to such an approach. The standard approach was to discuss 'whether the three JTB conditions were necessary' and then to explore sufficiency via the original Gettier cases. Students then tended to introduce the 'no false lemmas' response, dismiss this using the 'Barn County' example, before outlining and dismissing infallibilism for being 'too strict'. At this juncture, students tended to proceed in one of two ways – either by explaining and dismissing a version of simple reliabilism and concluding in favour of virtue epistemology, or by focusing solely on reliabilism, with better responses concluding that a more sophisticated version of this theory (typically Nozick's or Goldman's) does the best job of avoiding the pitfalls that other definitions face – eg 'by countering Gettier cases and allowing for children and animals to have knowledge'. Some argued that there was no successful definition of propositional knowledge, whilst others argued that

whilst there was no fully satisfactory definition, one was better than the others, with some arguing for infallibilism despite its restrictions. This latter approach can produce a perfectly valid argument, but students who took it often failed to give sufficient detail and so did not end up presenting a convincing case as to why infallibilism was the most successful definition.
Some good use was made of Zagzebski's views on the nature of definition, particularly when this was used to unpack the strengths and weaknesses of different definitions of propositional knowledge, although too often students referred to Zagzebski's criteria without making any use of them in their argument. Similarly, some made use of the distinction between the different types of definition – eg 'real vs. nominal' and 'theoretical vs. practical', but more often than not this material was introduced without it being integrated into the overall argument. Whilst students were cognisant of this material, many were unable to use it with any analytical or evaluative impact. As such, it tended to add very little to the overall argument, and whilst it was not entirely irrelevant to the question, it could have been dispensed with and the time better spent developing evaluations in more detail.

Where students performed well

Students who accessed the top two bands did so through a careful balancing of objections and counter-responses leading to well-supported judgments throughout and overall. These responses sustained a clear and coherent line of argument and explicitly attributed appropriate weight to arguments being considered. The very best answers offered a robust defence of the position they were advancing and often made clear reference to crucial arguments in the introduction before expanding on it later in their essay. Robustness is a key feature of the 21-25 band and students who tried to cover all the definitions on the *Specification* often ran out of time, steam, or evaluative points, and so their responses fell short of being sufficiently robust. However, if they had shown good evaluative skill on the majority of the content covered, then such responses could still be placed high in the 16-20 band.

Some good answers adopted the standard approach outlined above, covering a lot of ground, and did it well. This often involved an attempt to robustly defend one of the main accounts, showing how their favoured definition 'dealt with the Gettier and Barn County cases and so provided both the individually necessary and jointly sufficient conditions for knowledge', which they argued were 'an essential requirement for a good definition'. Some then engaged in additional evaluation – eg by explaining why children and animal knowledge works in favour of reliabilism over its competitors, or by defending virtue epistemology against 'criticisms of vagueness'. When it came to discussing children and animal knowledge, some good responses questioned whether such knowledge was 'more ability than propositional', which added an extra dimension to their evaluation and enabled them to integrate points made earlier in their essay regarding the three types of knowledge.

Some students who went for breadth dealt with 'issues concerning the necessity of the tripartite conditions' relatively quickly. This allowed for time to engage fully in evaluation of sufficiency problems and the full range of responses to this. Others went into more detail on necessity issues and were able to handle objections and responses to them well. However, it must be noted that the weakest part of such evaluations often pertained to the truth condition. The issues here are tricky and even better students struggle with 'Kuhnian relativism', often introducing some of Kuhn's ideas but not developing them precisely or making their impact for the concept of truth fully clear. Others took a different approach – eg by briefly outlining the necessity of the tripartite conditions to begin with and then going straight into sufficiency objections before circling back to the problematic nature of justification, and then developing this through an evaluative discussion of reliabilism and/or virtue epistemology.

Other responses were more selective, but equally effective (more so in some cases) in constructing complete and coherent arguments with a careful balancing of positions. These tended to focus on one theory from the start – eg reliabilism or virtue epistemology – and then defended it against a series of increasingly forceful objections. This approach allowed the student to select and integrate the strongest arguments and counters more effectively, as well as allowing for critical comparisons between their favoured definition and the alternatives. Students who adopted this approach were often able to demonstrate their breadth of understanding in a more analytical and evaluative way than those that attempted to cover all the definitions.

Common / noteworthy weaknesses in student responses

Some of the characteristics of the middle level responses (11-15 band) was the level of clarity, detail and precision of points raised, the depth of the arguments and counter-arguments provided, and the coherence of the overall line of argument. Whilst the overwhelming majority of students were able to give accurate or generally correct accounts (including some evaluation) of a good portion of the relevant material - eg the tripartite view, Gettier's original example/s, the no-false lemmas response, infallibilism and 'Henry in Barn County' - there was significant variation in terms of the levels of detail and precision demonstrated. Where students often came unstuck was in their discussions and evaluations of reliabilism and/or virtue epistemology. For example, students were generally good at describing virtue epistemology (using Sosa's archer analogy), but not so good at applying it to problematic cases such as Barn County. Some thought it could deal with it; some thought it could not; but few spent much time talking about what exactly intellectual virtue might mean in practice, which could have been an important part of their evaluative discussions. Reliabilism was another theory that was not discussed or evaluated particularly well at the lower end, with key points often being passed over too quickly. Little time was spent discussing what 'a reliable cognitive process' actually is and how this differs from a justification. This could have been approached in relation to the distinction between externalism and internalism making the analysis and evaluation more compelling. Lower down this band, many responses failed to give adequate consideration to possible responses to objections – eg different versions of reliabilism – that meant these answers often lacked depth and balance. Finally, while there was an improvement this year, there were still quite a large number of students who were asserting that a particular position is 'stronger' or 'weaker' without further explanation or justification, while using too many stock phrases such as 'this is a strong response' or 'this is a crucial argument' without ever explaining *why* it is strong/crucial.

Further down the performance scale (6-10 band), there were also weaknesses in the construction and coherence of the overall argument. At times, it appeared that some students had tried to rote learn an essay plan without fully grasping the material. Some got themselves muddled up when trying to recall the essay plan and ended up by wrapping themselves up in knots. An example of this would be when students argued that 'truth, belief, and justification were neither necessary nor sufficient conditions for knowledge', but then concluded that 'reliabilism' or 'no-false lemmas' was a successful definition, despite having argued against the necessity of conditions such as truth and belief. This was a feature of the 6-10 band where an argument was presented 'but not fully coherently'. Weaker responses also tended to move through the material too quickly, often listing definitions with limited integration or by giving lengthy descriptions of Gettier and Barn County cases without engaging in detailed evaluation. This involved juxtaposing and describing a range of views without sufficient engagement in ongoing evaluation or replacing argument with assertion.

Among students who did not make it out of the 1-5 band, they either tended to be able to say very little beyond making a few points about the different kinds of knowledge (including propositional), the JTB definition of knowledge, or else there was confusion with other epistemological topics on the *Specification* – eg taking the issue to be 'whether rationalists or empiricists give the correct definitions of knowledge'. Some credit was given to the latter where there were (even tangentially) relevant points raised.

Section B: Moral Philosophy

Question 6: State Ayer's verification principle. (3 marks)

General Observations

This question tested students' knowledge and understanding (AO1) of a central philosophical principle on the *Specification* - 'Meta-ethics' (3.2.3) - but which also appears in the 'Religious language' section (3.3.3) of the Metaphysics of God topic on Paper 2 (7172/2). As such, there was no requirement for students to explain the principle solely in relation to Ayer's meta-ethical views and, whilst a good number of students chose to do so, many explained the principle in isolation of any particular application.

Students performed better on this question than on the corresponding Epistemology 3-mark item, with just under 75% gaining 2 marks or above, and well over half of students achieving maximum marks. According to the mean average, this was the second best performing item on the paper after Question 02. But at the bottom end over 14% of students were not able to score any marks, and with the exception of Question 08 (see the discussion below) more students failed to even attempt an answer to this question than any other on the paper this year.

Where students performed well

There were 3 key aspects to this question: (1) the verification principle as a theory of meaning/literal significance; (2) the notion of empirical verifiability; and (3) the notion of the analytical (or tautological). To get the full 3 marks, students needed to address each aspect clearly and precisely and, whilst students did not need to use Ayer's exact phrasing/terminology, many of the best responses did. Some of these simply stated the verification principle without any supporting examples or elaboration - eg in terms of 'strong/weak verification' or 'verification in principle/in practice', etc. Other students gave a clear and correct account of the principle and then chose to develop this, either through examples, a brief explanation of key terms such as 'empirical verifiability' or 'analytic,' or through further developments of the principle. As with question 1, examples were not required, but they often added value to the responses.

Common / noteworthy weaknesses in student responses

There were a number of reasons why students failed to gain full marks. Answers that gained 2 marks tended to miss out a key element. Answers that gained 1 mark tended to be clear on one aspect only or demonstrated a partial or imprecise understanding of more than one aspect - in both cases, the answer was deemed to be fragmented. There are numerous examples listed in the indicative content of the Mark Scheme that exemplify a range of the types of responses seen.

One of the most common errors of understanding involved students not fully grasping that the verification principle is a criterion of *meaning*. A number of students appeared to conflate Ayer's principle with ‘Hume’s Fork’, stating that the verification principle is the claim that ‘something only counts as *knowledge* if it is either empirically verifiable or analytically true’. Whilst there is a clear relationship between Hume and logical positivism (indeed, at one point Ayer describes *Language, Truth and Logic* as ‘Hume in modern dress’) the failure to understand the verification principle as a criterion of meaning meant the concept was not fully understood, although students could still gain 2 marks if the rest of the answer was correct. Another error, although less common, was to conflate verification with falsification.

Question 7: **Explain Aristotle’s function argument. (5 marks)**

General Observations

This question required knowledge and understanding (AO1) of a key philosophical argument from the ‘Normative ethical theories’ section of the *Specification* (3.2.1). Students performed better on this than Question 08 but less well than on both Epistemology 5 markers. The mean average mark was approaching 60% of available marks, with nearly two thirds of students managing to score at least 3 marks, demonstrating that they had grasped the substantive content of what the question was assessing. Over a quarter of students scored 4 marks. In order to produce these high scoring answers, they needed to show some understanding of what Aristotle means by ‘function,’ what the unique human function is (ie reason) and how fulfilling this function constitutes living well/flourishing. Many used examples/analogies to illustrate the concept of a function (eg ‘the function of a knife’) and, when used well, these helped students to bring out the salient points of Aristotle’s argument (eg ‘a good knife is one that cuts well, and a good human life is one that is lived well – ie in accordance with a rational principle’). Just over 15% of students were awarded the maximum number of marks. At the bottom end, less than 8% of students were unable to any score marks, while almost 19% scored 2.

Where students performed well

At the upper end of the performance scale, students who grasped the substantive content of the argument were then able to go on to give fuller, clearer, and more precise accounts, which demonstrated a more sophisticated understanding of Aristotle’s argument and the key concepts involved. This typically involved more fulsome explanations of why ‘reason is the unique human function’ (as opposed to nutrition or perception, which is also common to non-humans) and how fulfilling this function well (ie virtuously or excellently) leads to *eudaimonia*. Two things that stood out in higher level responses were a close relationship between the answer and Aristotle’s argument as set out in the *Nichomachean Ethics* (see below) and the precise and effective use of Aristotelian vocabulary such as *ergon* (function), *telos* (purpose/end), *eudaimonia*, and *arête* (excellence). When used in the context of a full and correct answer, the use of such terms serves to add precision.

Common / noteworthy weaknesses in student responses

Responses that gained less than 3 marks lacked the substantive content, giving some relevant points but not precisely, and so not demonstrating a sufficient understanding of Aristotle’s argument. This occurred when students missed out key aspects of the argument – eg reason being the unique human function - or when they used an example (eg the knife) but failed to explain how it related to the human function. Other errors occurred when students focused on different aspects of Aristotle’s virtue ethics such as the ‘doctrine of the mean’ or when they gave some relevant points – eg about *eudaimonia* – but failed to give sufficient detail on the function argument. Some students got muddled with the idea of a function and thought that it meant ‘working together harmoniously’ – eg how a good society might function well if everyone works together.

Another common imprecision (or over simplification) worth mentioning is when students claimed that Aristotle argues that ‘Everything has a function, so humans have a function’. Some secondary sources and commentaries do suggest that he believes that ‘everything has a function’ and that this is part of his broader teleological worldview. However, Aristotle does not refer to this in his presentation of the function argument in the *Nichomachean Ethics* and for purposes of accuracy and consistency it is advisable to be guided by the original text when it comes to the named arguments of specific philosophers. As such, the claim that ‘everything has a function’, stated bluntly without the necessary nuance – ie ‘for all things *that* have a function, the good resides in that function’ – was often a feature of lower scoring answers. However, such a generalised claim about the scope of ‘function’ in Aristotle’s philosophy would always be taken within the context of the rest of the response and would not preclude a student from accessing the full range of marks if the rest of the response was clear, full, and precise in explaining in the details of the function argument.

Question 8: Explain Mackie’s argument from relativity against moral realism. (5 marks)

General Observations

This question required knowledge and understanding (AO1) of one of Mackie’s anti-realist arguments from the ‘Meta Ethics’ section of the *Specification* (3.2.3). On average this was the lowest performing 5-mark question on the paper, although the mean was slightly higher than the mean for the paper as a whole. There was a greater preponderance of low scores here: over a quarter of responses were awarded 1 or 2 marks; over 14% of students failed to register a score; and more students did not attempt to answer it than any other on the paper this year.

Despite this, there were some excellent responses, and over 60% of students did manage to score at least 3 marks, showing that they had grasped the substantive content of the argument. In order to do this, they needed to demonstrate an understanding of the central claim of moral realism – ie concerning the existence of mind-independent moral facts - and why the variation and disagreement in moral codes are best explained in anti-realist rather than realist terms, rendering moral realism less likely to be true than moral anti-realism. Students often gave their favourite examples of moral relativity. The better performing responses did so by explicitly linking the example to the argument, whilst weaker responses tended to give lurid and elaborate descriptions without fully explaining why the proliferation of such cases presents a particular problem for the moral realist. In effect, many weaker responses simply stated that ‘there can't be any moral facts because otherwise everyone would agree about what is right and wrong’.

Where students performed well

At the upper end of the performance scale, students who grasped the substantive content then went on to give fuller accounts, which demonstrated a more precise and detailed understanding of Mackie’s argument. This typically involved students explaining the abductive nature of Mackie’s argument, an explanation of the two rival hypotheses in play, and why the anti-realist hypothesis is the better of the two.

Common / noteworthy weaknesses in student responses

As with Question 03 on Russell’s ‘best hypothesis’, it was surprising that more students did not refer to the abductive nature of Mackie’s argument. An ability to employ key philosophical vocabulary, including the use of logical terms, is one way in which students can add precision to their answers. Those students that did make use of such terms tended to perform better. Weaker responses missed out the abductive aspect and stated – imprecisely - something along the lines of ‘the fact there are differences in moral beliefs shows that moral realism is false’.

Responses that gained 2 marks lacked the substantive content, giving some relevant points but not precisely, and so not demonstrating a sufficient understanding of Mackie’s argument. For example, some students knew that Mackie was ‘a cognitivist anti-realist’ who advocates an ‘error theory’ of moral discourse, and they were able to explain/illustrate moral relativity, but they failed to link these points together, and so the logic of Mackie’s argument was unclear or just not present. Other common errors included students either mixing up Mackie’s arguments by presenting hybrid versions of the argument from relativity and the argument from queerness, or by straightforwardly presenting the argument from queerness rather than relativity.

As with Question 03 on Russell, Ockham’s razor was often referred to. While this is not an obvious misuse of Ockham’s razor given that moral anti-realism is more ontologically parsimonious than moral realism, it is still not something that Mackie refers to in his presentation of the argument in *Ethics: Inventing Right and Wrong*. Students were not penalised for mentioning Ockham’s razor if the rest of their answer was full and precise. But in some cases, students relied too much on Ockham’s razor and, as with Russell, seemed to confuse ontological simplicity with explanatory simplicity by claiming that ‘Mackie’s theory is simpler’ but without explaining in what way(s). As noted above on the Aristotle question, for purposes of accuracy and consistency it is advisable to be guided by the original text when it comes to the named arguments of specific philosophers.

Question 9: Explain how Kantian deontological ethics might be applied to the issue of simulated killing. (12 marks)

General Observations

This question required knowledge and understanding (AO1) of how a key normative ethical theory might be applied to the issue of simulated killing from the ‘Applied ethics’ section of the *Specification* (3.2.2). The question was well answered with the mean average above that of the paper as a whole, and students performing roughly 0.5 marks better than on the corresponding Epistemology question. Over half the students scored a mark of 7 or higher, with just over 20% scoring a mark in the top band (10-12). Over a quarter of students were awarded a mark in the 4-6 band, with less than 14% of students placed in the bottom scoring band (1-3). Less than 3% of students failed to pick up any marks at all.

To access 7 marks or above on this question, it was essential for students to provide relevant and correct material relating to both parts of the question. A response that only explained Kant’s deontological position but did not clearly address how his approach might apply to the issue of simulated killing could not demonstrate the level of detailed understanding and integration required for marks higher up the scale. Such an answer could, at best, be described as ‘largely correct’ with ‘most points made clearly’. Similar to the Epistemology 12 marker, where students only produced a partial response – omitting or otherwise failing to score marks on the second part of the question - they could still gain up to 6 marks for a precise and detailed explanation of Kant’s deontological ethics. In reality, however, where students demonstrated good knowledge and understanding of Kant’s deontological approach, some success on the second part of the question was pretty much a given: the majority of these students had *some* knowledge of what constitutes simulated killing and were therefore able to integrate this with aspects of Kantian moral philosophy.

As with the Epistemology 12 marker, many students approached this question in two parts. They would begin by explaining Kant's deontological ethics, including the two formulations of the categorical imperative, and then apply this to the issue. Many students acknowledged that there were different ways of interpreting how Kant would have responded to the issue, and some were able to unpack their knowledge of Kant effectively by drawing out these different aspects. Occasionally, students would bring in other ethical theories as a point of comparison, which often constituted redundancy, but overall students tended to show their adeptness in meeting the assessment objectives in relation to the relevant material.

Where students performed well

Most students tended to apply both the first and second formulations of the categorical imperative to the issue although there was no requirement to do so. Equally, if they did apply both, there was no requirement to give equal time and space to each, and some successful responses focused on one formulation more than the other. Some students managed to give very detailed and nuanced explanations of just one formulation and how it could be used to provide a response to the issue of simulated killing – eg by showing that 'whilst simulated killing does not directly contravene the humanity formula given the lack of rational agency, if it did result in participants developing negative traits that caused them to not treat others as ends-in-themselves, then it could be argued that Kantian's would disapprove'. They might then go on to acknowledge that a lack of empirical support to establish such a causal connection would also need to be taken into consideration. This was typical of more sophisticated responses, which tended to acknowledge that while 'simulated killing can pass the test of the categorical imperatives, Kant would question the effect it has on rationality and autonomy'. Another key feature of answers in the top band was the way in which Kantian terminology was precisely used and the manner in which the logic of Kant's approach was set out in precise detail. One final thing to mention is that some of the best responses integrated Kant's deontology and how it applies to the issue right from the outset – eg by showing a precise understanding of some of the ethical issues arising from simulated killing, outlining the first formulation and applying, and then outlining the second formulation and applying that.

Common / noteworthy weaknesses in student responses

Those who performed less well did so for a number of reasons. Weaker responses tended to be imprecise on both Kantian deontology and the applied issue or failed to address both parts of the question - with some only explaining Kant's deontology and others only discussing simulated killing, often using lots of descriptive examples from computer games but without addressing the key issues. There was some confusion between *simulated* killing and *actual* killing, with some students stating that 'in simulated killing people are being used merely as a means for my entertainment', which is obviously not accurate. Whilst many seemed to know Kantian ethics reasonably well, there was one blind spot that was often evident. Many were not fully clear that for Kant, unlike utilitarians, most actions are morally neutral. Too many assumed that if we do not have a duty to perform an action it must be prohibited. This meant that some ended up with the view that 'because simulated killing is not morally obligatory, then it must be morally wrong'; or alternatively, that 'because we can universalize simulated killing then we all have a moral obligation to practice it'.

Question 10: Can utilitarianism be successfully defended? (25 marks)

General Observations

This was the second of two questions on the paper designed to test both AO1 and AO2, but mainly A02. The question centred on assessing the extent to which utilitarianism - one of three main theories from the topic ‘Normative ethical theories’ (*Specification* 3.1.2) - can be defended. With students on average picking up just over 47% of the available marks, this was the lowest-scoring question on the paper, although only marginally, with Question 05 performing slightly better. However, the mean average - fractionally short of 12 - was higher than on the corresponding question in 2023 (on Aristotelian virtue ethics), when it was 11.05. This is largely accounted for by stronger performances among students in the middle range. But more students also accessed the top two scoring bands, with over 21% of students awarded at least 16 marks, while the number of students awarded marks in the top band (21-25) also increased and is currently approaching 5%. The most frequently accessed level was 11-15, with approximately 40% of students awarded a mark in this band – 13 was the typical score. Lower down, just over 30% of responses were placed in the 6-10 range of marks, while less than 9% of students scored 1-5 marks. It was striking that less than 1% of students failed to pick up any marks on this item of assessment: almost every student was able to write something relevant, even if just a sketchy explanation of the trolley problem with a cursory nod to a crude version of utilitarianism.

As with Question 05, answers were often very similar in structure and involved students working through the main objections and responses by incorporating different versions of utilitarianism. The standard approach was to discuss act utilitarianism, rule utilitarianism, then preference utilitarianism, with each theory being seen to build on some of the weaknesses/issues of the previous theory with varying degrees of success. Some mentioned Moore’s ‘ideal utilitarianism’ but given that this is not on the *Specification* this was understandably less common and not expected. In a majority of cases, the overall conclusion was that utilitarianism was not an effective moral theory, often due to the issue of moral integrity, which was generally argued to be the ‘crucial’ and most damning issue as it cuts across the three main types – unlike calculation, tyranny of the majority and pleasure not being the only good, which can arguably be dealt with by either rule or preference utilitarianism.

As expected, as well as sticking to the three main types of utilitarianism, the vast majority of students stuck to the main objections from the *Specification* and illustrated these with tried and tested examples - eg ‘framing an innocent person’, ‘throwing Christians to the lions’, ‘saving the cancer doctor over a loved one’, etc. The ‘mad axeman’ make his regular guest appearance, this time ruining the party for ‘strong rule utilitarians’. Other students discussed meta-ethical objections such as Moore’s ‘open-question argument,’ the ‘Naturalistic Fallacy’ and Hume’s ‘is/ought gap’. When the latter was done well - eg by linking to Mill’s ‘proof’ - this added precision and robustness, which is a typical feature of answers in the top band. However, some students got quite muddled with the application of these objections, which needed to be better understood before being utilised effectively in the exam.

The vast majority of students started with Bentham and gave descriptions of the felicific calculus. In better responses this was well integrated into problems regarding calculation, although weaker responses tended to spend too long on AO1 - detailing Bentham’s version of ‘act utilitarianism’ - rather than getting down to evaluation. Many did go on to discuss problems with Bentham’s quantitative hedonistic act utilitarianism, such as the ‘quality of pleasure’ or the ‘tyranny of the majority,’ before introducing Mill as a ‘qualitative rule utilitarian’, as a way of responding to these objections. From here, many students opted to introduce Nozick’s ‘experience machine’ as an objection to both Bentham’s and Mill’s hedonistic utilitarianism, before turning to preference utilitarianism as a response. Treatments of

Nozick varied. Whilst many used his thought experiment to good effect, quite a few thought that utilitarians would have to force everyone into the experience machine (sometimes ‘at gunpoint’). Others tried to use Mill’s distinction between higher and lower pleasures as a response to Nozick but failed to realise that Nozick’s thought experiment is set up to accommodate this distinction. In some cases, students were aware of this and set up two versions – one where ‘the machine gives only lower pleasures’ and a second version where ‘it gives both lower and higher’ – then used Mill to combat the first whilst showing how his distinction fails to address the second. This worked well in some cases but less so in others and not at all in those cases where students thought that Mill’s distinction can be used to respond successfully to Nozick.

Continuing on the theme of Mill’s distinction between higher and lower pleasures, in many cases the standard critical response was that Mill’s distinction is ‘elitist’, with little further explanation as to why this is a problem. Is the argument that there is actually no real distinction between higher and lower pleasures, and that it is all, therefore, just a matter of Victorian class-based snobbery? If this is the case, then more was needed on why exactly the distinction is merely a reflection of social class and why that is an issue. Alternatively, is the point that there is a distinction, but because only the elite have ready access to higher pleasures, then this is morally objectionable in some way? On the other hand, it could be argued that maximising higher pleasures could be the opposite to elitism – eg a society that maximises higher pleasures is one where arts education is, for example, spread throughout the education system, which could be regarded as a positive good. More generally there was at times far too many bald assertions, with students taking critical points at face value as ‘knock down’ arguments and failing to develop or analyse them in a satisfactory way.

Where students performed well

Students who accessed the top two bands did so through a careful balancing of objections and defences leading to well-supported judgments throughout and overall. The best responses sustained a clear and coherent line of argument and explicitly attributed appropriate weight to arguments being considered. The very best answers offered a robust defence of the position they were advancing and often made clear reference to a crucial argument in the introduction before expanding on it in the final part of their essay and/or conclusion. Robustness is a key feature of the 21-25 band and students who were able to produce a strong defence of a position – eg a particular version of utilitarianism or arguing against the theory as whole – were more likely to be awarded marks in this band accordingly, although all the level descriptors are taken into consideration.

As mentioned above, some students used meta ethics in their responses very well indeed - eg arguing that ‘utilitarianism fails because naturalism fails’ - or alternatively defending utilitarianism against such objections. Another feature of higher-level responses was how well they dealt with preference utilitarianism and how well they were able to develop their evaluations on the ‘collapse of rule into act utilitarianism’. Given the wealth of material available, better responses tended to stick to utilitarianism although some students effectively conducted some of their evaluation in relation to other normative theories whilst being very careful not to juxtapose, stray into irrelevance or get bogged down in a general discussion of Kant or Aristotle (which was often the opposite with answers in the lower bands).

Common / noteworthy weaknesses in student responses

Where students scored marks in the bottom level (1-5) they tended to produce either very short responses with substantial gaps in the content, or a longer response that had very little/no evaluation, no clear line of argument and/or where there was serious misunderstanding.

Responses that scored in the 6-10 band tended to have a generally sound degree of knowledge and understanding, which was often shown through lengthy descriptions of key points – eg of the felicific calculus or higher/lower pleasures - or through unnecessarily long retellings of various classic thought experiments. However, such responses were not entirely devoid of AO2, and they would go on to raise some relevant critical points in a mostly coherent structure. A typical feature of the 6-10 band is arguments ‘not being fully coherent’. It was not uncommon to see some students struggling to construct a coherent argument by not carefully distinguishing between different types of utilitarianism and/or failing to integrate arguments and positions – eg by arguing against hedonistic utilitarianism, dismissing it, then discussing preference, and dismissing that, before returning to hedonistic (perhaps in the form of a rule version) without seeming to realise that they had already ‘proven’ hedonistic utilitarianism to be ‘false’. Weaker answers also tended to juxtapose criticisms rather than critically engaging with them and so the logical linking between objections and defences was not always clear. Overall, answers in this band were typified by a lack of detail and precision in the points raised, as well as limited meaningful engagement with the issues.

As already mentioned, one version that tended not to be done as well was preference utilitarianism, despite the fact that many included some discussion of it in their response. Most students can explain clearly act and rule utilitarianism, but there was too much vagueness on preference utilitarianism. This might in part be because it *just is* a vaguer theory and, if so, this could be used as a relevant critical point if suitably developed. Quite a few did not seem to know what ‘preference’ means in the context of utilitarianism, and many were hazy about exactly whose preferences might count in any given moral dilemma. An example of this is when students discussing the ‘Jim and the Indians’ scenario suggest that ‘Jim’s preference not to kill anyone’ is the only preference that might matter in such a situation.

As with Question 05, some essays ended up being a list of points – one after the other without integration and without development. There are still a number of students asserting that a particular objection or defence is ‘stronger’ or ‘weaker’ without explanation. And again - too many stock uses of phrases such as ‘this is a strong response’ or ‘this is a weak argument’ without ever explaining why it is weak/strong with evidence. But this is less intrusive than it has been in previous series, and more students are aligning their statements about what they are doing with the actual arguments they are producing.

Mark Ranges and Award of Grades

Grade boundaries and cumulative percentage grades are available on the Results Statistics page of the AQA Website.