AI & CAREER DEVELOPMENT · POSITION ARTICLE

When the Bot Is Free, Your Students Are the Product

What We Need to Know About For-Profit AI Career Coaching

David Carney · Executive Director, Career Industry Council of Australia

There is a version of this story that sounds straightforwardly positive. Artificial intelligence is coming to career guidance in Australian secondary schools. Thousands of students will gain access to a free AI career coaching tool. Young people who have never had access to a qualified career practitioner will finally get some form of support. What is there to object to?

Quite a lot, as it turns out. Not because AI has no role in career development, it does, and the profession should engage with that seriously, but because the way these initiatives are typically structured raises questions that deserve honest answers before we hand over our students, their data, and their futures to commercial platforms.

This article is not a rejection of technology. It is a call for the kind of rigorous, honest scrutiny that any service reaching vulnerable young people at a critical point in their lives should expect as a matter of course. And it is a call to the profession, to school leaders, to parents, and to government to ask harder questions than the announcement invites us to ask.

The Business Model Behind the Bot

The first thing to understand about any free AI tool deployed at scale is that free is a pricing decision, not a charitable one. Technology companies do not absorb the cost of building, training, and deploying AI systems out of goodwill. They recover it, typically through the data generated by users, the commercial relationships built around the platform, or both.

This matters enormously when a commercial platform is given AI-powered access to the career aspirations, anxieties, and personal data of secondary school students. Every interaction a student has with an AI career coaching tool, every question asked, every pathway explored, every piece of data entered, has potential commercial value. It trains the model. It refines the product. It builds a dataset of considerable worth. The students using these tools are not the customers. They are the infrastructure.

It is also worth examining the broader commercial ecosystems in which some of these platforms operate. Platforms that deliver brand-sponsored content programs alongside career guidance are not operating a neutral service. The commercial model is not incidental to the product. It shapes it. When those same

platforms are positioned as providers of AI career coaching, the question of whose interests the tool is designed to serve becomes unavoidable.

None of this makes free AI career tools automatically harmful. But it makes transparency and governance non-negotiable, and it makes the question of consent more complex than it first appears.

In a school context, students are minors. In most cases, they will not be the ones providing consent to participate in an AI career coaching program. Their school will, typically, on their behalf, often through a terms of service agreement that receives little scrutiny.

That is a meaningful ethical distinction. A principal signing up to a free AI platform on behalf of hundreds of students is making a consequential decision about those students' data, their exposure to commercial influence, and the quality of guidance they will receive. They deserve to make that decision with full information, not on the basis of a press release and a free trial.

Principals, parents, and students have a right to know, before engaging with any AI career tool, exactly what data is being collected, how it will be used, who it will be shared with, and whether it will be used to train or improve a commercial product. Informed consent is not a bureaucratic nicety in this context. It is an ethical obligation.

The Engineered Influence Problem

There is a subtler risk that sits alongside the data question, and it is in some ways more concerning because it is harder to see.

Career guidance involves values. When a qualified career development practitioner works with a young person, they bring professional ethics, a duty of care, and a commitment to the client's interests above all else. They are bound by professional standards. They can be held accountable. The guidance they offer is, at its best, genuinely impartial.

An AI system encodes values whether its designers acknowledge this or not. Every weighting, every recommendation pathway, every way the system frames options and presents information reflects choices that were made by someone. The question is whether those choices were made in the interests of the student, or in the interests of the platform and its commercial partners.

CICA's Guiding Principles for Career Development Services and Career Information Products are explicit on this point. Career guidance must be impartial. It must present options honestly, including options that do not serve the provider's interests. It must be designed around the client, not the business model.

An AI career tool that steers students toward partner institutions, that weights certain pathways because a sponsor has paid for visibility, or that presents options more prominently because they serve a commercial relationship, is not career guidance. It is engineered influence. And because it operates at scale, through an interface that feels neutral and personalised, it has the potential to shape the career trajectories of an entire generation in ways that serve commerce rather than young people.

Without independent auditing of the values and weightings embedded in these tools, we have no way of knowing whether this is happening. That absence of transparency is itself a governance failure, and it is a failure that Australia's current regulatory environment is not equipped to address. Unlike some comparable jurisdictions, Australia has no specific regulatory framework governing the use of AI in educational settings.

There is no mandatory disclosure requirement, no independent auditing mechanism, and no clear accountability pathway when an AI system causes harm to a student. That regulatory vacuum is not a minor oversight. It is the environment in which these tools are currently being deployed.

It is also worth acknowledging a tension that this discussion would be incomplete without naming. When AI career coaching initiatives are government-backed, government is not a neutral regulator standing outside this question, it is a party to it. Calling on government to audit and govern an initiative it has publicly endorsed requires a degree of political will that cannot be assumed. That is precisely why the profession, school leaders, parent communities, student advocates, and other stakeholders need to be actively engaged in this conversation, not waiting for a government response that may not come unprompted.

Bias, Hallucinations, and the Danger of Confident Misinformation

Commercial influence is a risk we can at least conceptualise and guard against. There are two further technical properties of AI systems that are less well understood outside the profession, and which have serious implications for career guidance specifically.

The first is bias. AI systems learn from data, and the data they learn from reflects the world as it has been, not as it should be. When an AI career tool is trained on historical labour market data, graduate outcomes, and career pathway information, it absorbs the patterns embedded in that data, including patterns of gender segregation, racial and socioeconomic inequality, and structural disadvantage. A tool trained on decades of workforce data will, without deliberate correction, tend to reflect those patterns back at students. It may subtly steer girls away from engineering, underrepresent certain careers for students from lower socioeconomic backgrounds, or present pathways that reflect historical norms rather than contemporary possibilities.

But the equity harm is most acute where these factors intersect. A young woman from a low-income background in a regional or remote community does not face three separate bias risks. She faces a compounded one, where gender, socioeconomic disadvantage, and geographic isolation interact in ways that are qualitatively different from any single factor in isolation. These are precisely the students who have historically had the least access to quality career guidance, and they are the students most likely to be reached by a free AI tool deployed at scale. The promise of access is real. So is the risk that the access delivers guidance shaped by the very inequalities it should be helping to overcome.

This is not hypothetical. Amazon's internal AI recruitment tool, developed and ultimately abandoned after it was found to systematically disadvantage female candidates, is perhaps the most widely cited example of AI bias causing real harm, but it is far from isolated. Research on AI course recommendation systems, career assessment platforms, and automated advising tools has consistently found that bias in training data produces biased outputs, often in ways that are invisible to the user and deeply difficult to detect without rigorous independent testing. The student receiving a career recommendation has no way of knowing whether that recommendation reflects their individual potential or the accumulated weight of historical inequality. The interface feels personalised. The bias is systemic.

The second technical risk is hallucination. The tendency of AI language models to generate confident, fluent, and completely incorrect information. AI systems do not know what they do not know. When asked about entry requirements for a particular course, employment prospects in a specific field, salary ranges, or the existence of particular scholarships or programs, an AI career tool may produce detailed,

authoritative-sounding responses that are factually wrong. It will not flag its uncertainty. It will not say it is not sure. It will answer with the same confident tone it uses when it is correct.

In a career guidance context, the consequences of this can be serious and lasting. A student who receives incorrect information about entry requirements may not prepare appropriately. A student given inaccurate labour market information may make pathway decisions based on a version of the world that does not exist. A student incorrectly advised about the availability of financial support may miss an opportunity that could have changed the course of their education. These are not minor errors. They are the kinds of mistakes that affect real outcomes for real young people, and they may not become apparent until significant time and opportunity have passed.

Qualified career development practitioners make mistakes too. But they are professionally accountable for those mistakes, they can be corrected, and they carry an obligation to keep their knowledge current and to acknowledge the limits of what they know. An AI system has none of those qualities. This raises a question the current regulatory environment has not answered: if a student receives harmful or incorrect guidance from an AI career tool, what is the redress pathway? Whom do they complain to? What liability does the platform carry? In the absence of a specific regulatory framework for AI in educational settings, these questions have no clear answer. That is not an acceptable situation for a service being deployed to hundreds of thousands of young Australians.

The Mental Health Dimension

Career uncertainty and transition are significant contributors to youth anxiety and mental health distress. Research consistently shows that young people navigating educational and career decisions experience elevated levels of stress, and that the quality of support they receive during these transitions has measurable impacts on their wellbeing. Career development practitioners understand this. It is part of their training, part of their professional responsibility, and part of what distinguishes a qualified practitioner from an information service.

A conversation that begins with questions about career options can, in the hands of a skilled practitioner, become a moment of genuine support for a young person who is struggling. The practitioner notices the tone, reads between the lines, asks a different question, and responds to what is actually happening rather than what was asked. This is not incidental to good career practice. It is central to it.

An AI system cannot do this. It can be programmed to detect certain keywords and respond with crisis resources. But it cannot exercise the kind of human judgement that recognises distress in its early and ambiguous forms, before a student has found the words for it. For students who are already vulnerable, and for students most likely to be reached by a free AI tool in an under-resourced school, this is not a theoretical gap. It is a real one, with real consequences. Deploying AI career coaching tools without clear protocols for identifying and responding to student distress, and without guaranteed access to a qualified human when it matters, is a welfare risk that has received almost no attention in public discussion of these initiatives.

The Substitution Risk

Career counselling in Australian schools has been chronically underfunded. Career practitioners, where they exist, are frequently part-time, under-resourced, and stretched across responsibilities that go well beyond career guidance. Some jurisdictions have invested in training qualified human practitioners. Others have outsourced career education to commercial platforms. The results of those two approaches are visible.

The risk with the current wave of AI career tools is not simply that imperfect technology will be deployed. It is that free AI tools will be used to justify not investing in the qualified human workforce that schools actually need. The schools receiving an AI career coach will be the ones that never receive a qualified practitioner, because governments and school systems will conclude that the problem has been solved.

It has not been solved. Career development is a professional practice grounded in the Australian Blueprint for Career Development, the nationally recognised framework that articulates what quality career development looks like across the lifespan. The Blueprint, alongside the Professional Standards for Australian Career Development Practitioners and the Guiding Principles for Career Development Services and Career Information Products, represents decades of professional knowledge about what young people need to navigate their careers well. None of that knowledge is encoded in a commercial AI product. It requires trained, accountable people who can build relationships, exercise judgement, and respond to the full complexity of a young person's situation.

This is a point that students themselves deserve to hear directly. Young people engaging with AI career tools have a right to know that they are interacting with a commercial product, not a qualified professional. They have a right to know what data is being collected and how it will be used. They have a right to access a qualified human practitioner when their needs exceed what a digital tool can responsibly provide. And they have a right to expect that the guidance they receive is designed around their interests, not someone else's business model.

The Counter-Argument and Why It Matters

Any honest treatment of this issue needs to engage with the strongest case on the other side, because it is not a weak one.

Qualified career practitioners are also subject to bias. They also sometimes give incorrect information. They are also unavailable to the majority of Australian students, particularly those in regional and remote areas, those in under-resourced schools, and those whose schools have determined that career guidance is a lower priority than other demands. If the alternative to an AI career tool is no career guidance at all, which, for many students, is precisely the alternative, then the case for imperfect AI is stronger than critics of the technology may be comfortable acknowledging.

This argument deserves a direct answer. The answer is not that AI tools are worse than qualified practitioners across all dimensions and should therefore be rejected. The answer is that the choice being presented, AI tool or nothing, is a false one, created by decades of underinvestment that governments and school systems are now using AI to avoid addressing. Accepting that framing concedes the very argument we should be contesting. The profession's position should not be that AI is worse than a human. It should be that Australian students deserve both, properly funded human practitioners supported by well-designed technology, and that settling for less is a policy failure, not an inevitability.

What Good Looks Like

The career development profession is not opposed to technology. Used well, AI tools can genuinely extend the reach of qualified practitioners, support students in rural and regional areas who have historically had little access to career guidance, help young people explore options at their own pace, and free up practitioner time for the complex, relational work that only humans can do.

In practice, this might look like an AI tool that handles initial career exploration and information queries, flags students who need deeper support, and connects them to a qualified practitioner for follow-up. It might look like a tool that helps a student in a remote school access labour market information and course pathway details, with a practitioner available via video call when the conversation requires human judgement. It might look like AI-assisted career portfolios that help students track and articulate their developing skills and interests over time, with practitioner oversight built into the process. Technology in service of the profession, not in place of it.

Other countries are beginning to chart this path. In the United Kingdom, the Gatsby Charitable Foundation's Benchmarks for Good Career Guidance have been used as a framework for evaluating digital career tools, providing a model for how professional standards and technology governance can be brought together. Canada's career development sector has engaged seriously with questions of AI ethics and data governance in career services. These are not perfect models, but they demonstrate that it is possible to set standards, require accountability, and bring the profession into the governance conversation rather than bypassing it.

Getting there in Australia requires that AI tools deployed in schools demonstrate they are genuinely client-centred and have no commercial relationships that compromise the impartiality of their guidance. They must be independently tested and audited for bias, with particular attention to gender, socioeconomic background, cultural factors, and their compounding intersections. They must have transparent safeguards against the generation of inaccurate information, with clear escalation pathways to qualified practitioners. They must have explicit welfare protocols for identifying and responding to student distress. They must be subject to robust data governance, with explicit informed consent and a prohibition on the use of student data to train commercial AI products. They must be benchmarked against the Australian Blueprint for Career Development, the CICA Professional Standards, and the CICA Guiding Principles. And they must be evaluated on student outcomes, not reach and engagement metrics.

What School Leaders and Practitioners Should Ask Right Now

Systemic change takes time. While regulatory, professional, and policy frameworks catch up, school leaders and career practitioners are making decisions today about whether and how to engage with AI career coaching tools. Here are the questions that should be asked before any school signs on to any AI career guidance initiative.

- **Who owns this platform and what is their business model?** If the answer involves commercial partnerships, brand-sponsored content, or data monetisation, the next question is how those interests are structurally separated from the career guidance being delivered.
- **What data is collected, how is it stored, and who has access to it?** Read the terms of service, not the marketing materials. If student data is used to train or improve a commercial AI product, that should require explicit informed consent from families, not be buried in fine print.

- **Has the tool been independently tested for bias?** Ask specifically about gender, socioeconomic background, cultural background, and geographic factors. Ask how often testing is conducted and by whom. If the company conducts its own testing, that is not independent.
- **What happens when the tool gives wrong information?** Every AI system will generate incorrect information at some point. The question is what the safeguard is, how errors are identified, corrected, and communicated to affected students.
- **What are the welfare protocols?** If a student's interaction indicates distress or vulnerability, what happens? Is there a qualified human in the pathway, and how quickly can they be reached?
- **Is there a qualified career practitioner involved in the delivery of this program?** If the AI tool is being positioned as a replacement for a practitioner rather than a support for one, that is the wrong model.

These are not unreasonable questions. Any platform that cannot answer them clearly and transparently should not be in Australian schools.

The Question Worth Asking — And What CICA Is Doing?

CICA has issued a media release calling on the Federal Government to answer the questions left unanswered by the current wave of AI career-coaching announcements. We are asking for independent auditing of embedded values and commercial weightings, robust data governance frameworks that protect students, a regulatory framework for AI in educational settings, and a genuine, funded commitment to a qualified career development workforce.

We are also calling for a broader coalition. Student associations, teacher unions, parent organisations, university admissions bodies, and consumer advocacy groups all have legitimate interests in how AI career tools are designed, governed, and deployed. CICA does not claim to speak for all of those interests. But we are committed to working alongside them, because the questions raised by commercial AI in career services are too consequential to be left to any single professional body, or to the market.

> *When a for-profit company offers to provide a free service to schools, the first question is not how do we roll this out. It is what is this company getting in return, and is that exchange genuinely in the interests of students.*

Young Australians making high-stakes decisions about their futures deserve guidance that is honest, impartial, technically sound, free from commercial influence, and accountable to their interests alone. Whether an AI tool can provide that depends entirely on who built it, what data it was trained on, what values it encodes, how its errors are identified and corrected, and who is holding it to account.

Those are not questions the market will answer on its own. They are questions for the profession, for regulators, for school leaders, for parents, and for governments. And they need answers before we proceed.

SUGGESTED CITATION

Carney, D. (2026). *When the Bot Is Free, Your Students Are the Product: What We Need to Know About For-Profit AI Career Coaching.* Career Industry Council of Australia. https://cica.org.au

ABOUT THE AUTHOR

David Carney is Executive Director of the Career Industry Council of Australia (CICA), the national peak body for the career development profession. CICA maintains the Professional Standards for Australian Career Development Practitioners, the Guiding Principles for Career Development Services and Career Information Products, and the National Register of Professional Career Development Practitioners (RPCDP). To engage with CICA visit www.cica.org.au